Welcome To

ISSCC 2014
Session 1

Plenary

Computing's Energy Problem:

(and what we can do about it)

Mark Horowitz

Stanford University

horowitz@ee.stanford.edu

Everything Has A Computer Inside

The Reason is Simple: Moore's Law Made Gates Cheap

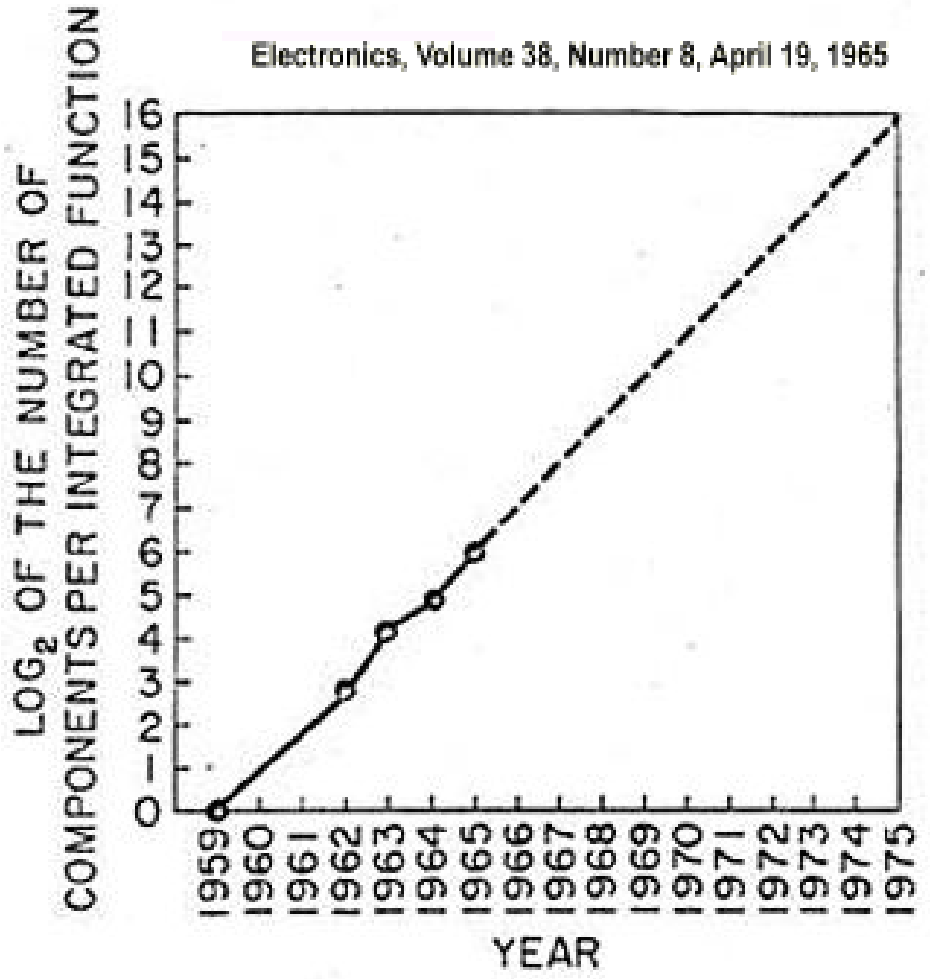

Fig. 2 Number of components per integrated function for minimum cost per component extrapolated vs time.

Dennard's Scaling Made Them Fast & Low Energy

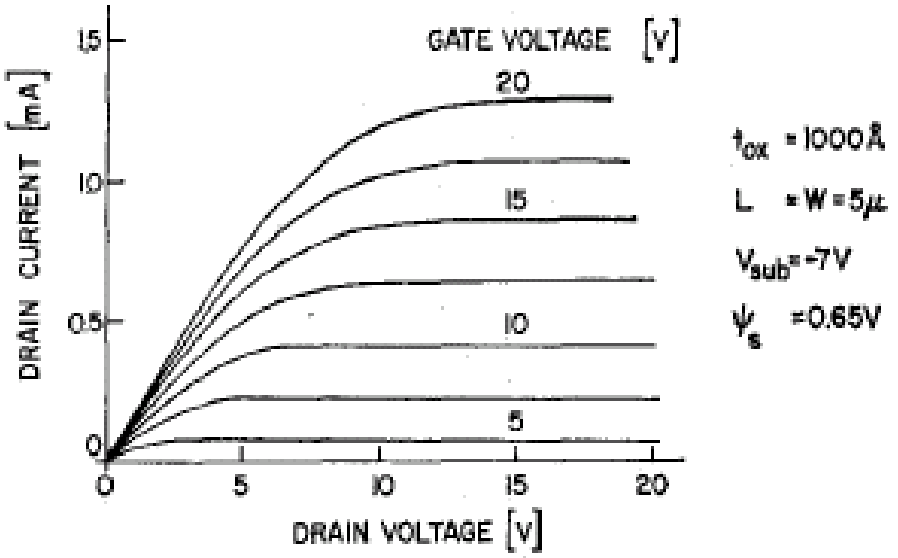

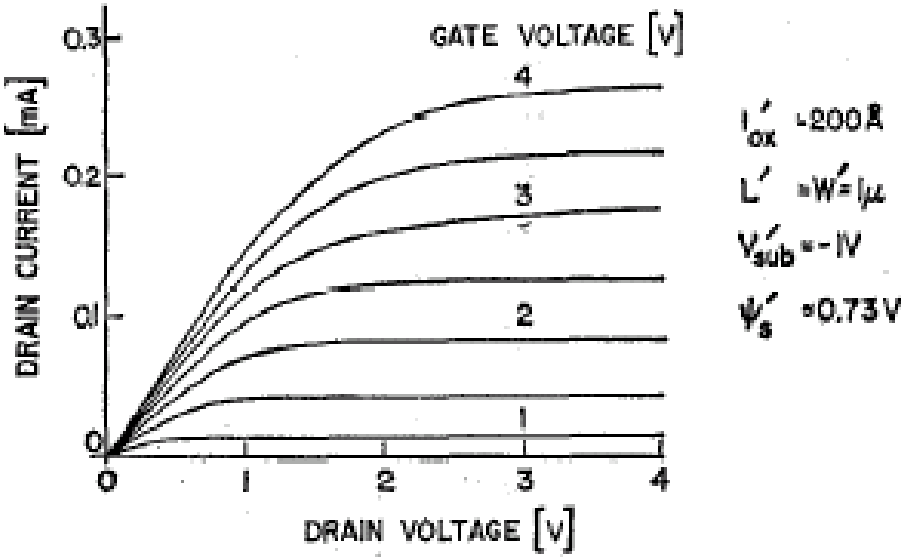

The triple play:

• Get more gates,	$1/L^2$	$1/\alpha^2$
• Gates get faster,	CV/i	α
• Energy per switch	CV^2	α^3

Dennard, JSSC, pp. 256-268, Oct. 1974

Our Expectation

Cray-1: world's fastest computer 1976-1982

- 64Mb memory (50ns cycle time)
- 40Kb register (6ns cycle time)
- ~1 million gates (4/5 input NAND)
- 80MHz clock
- 115kW

In 45nm (30 years later)

- < 3 mm^2
- > 1 GHz
- ~ 1 W

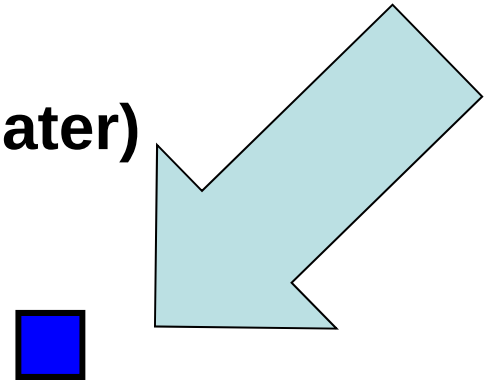

CRAY-1

Supporting Evidence

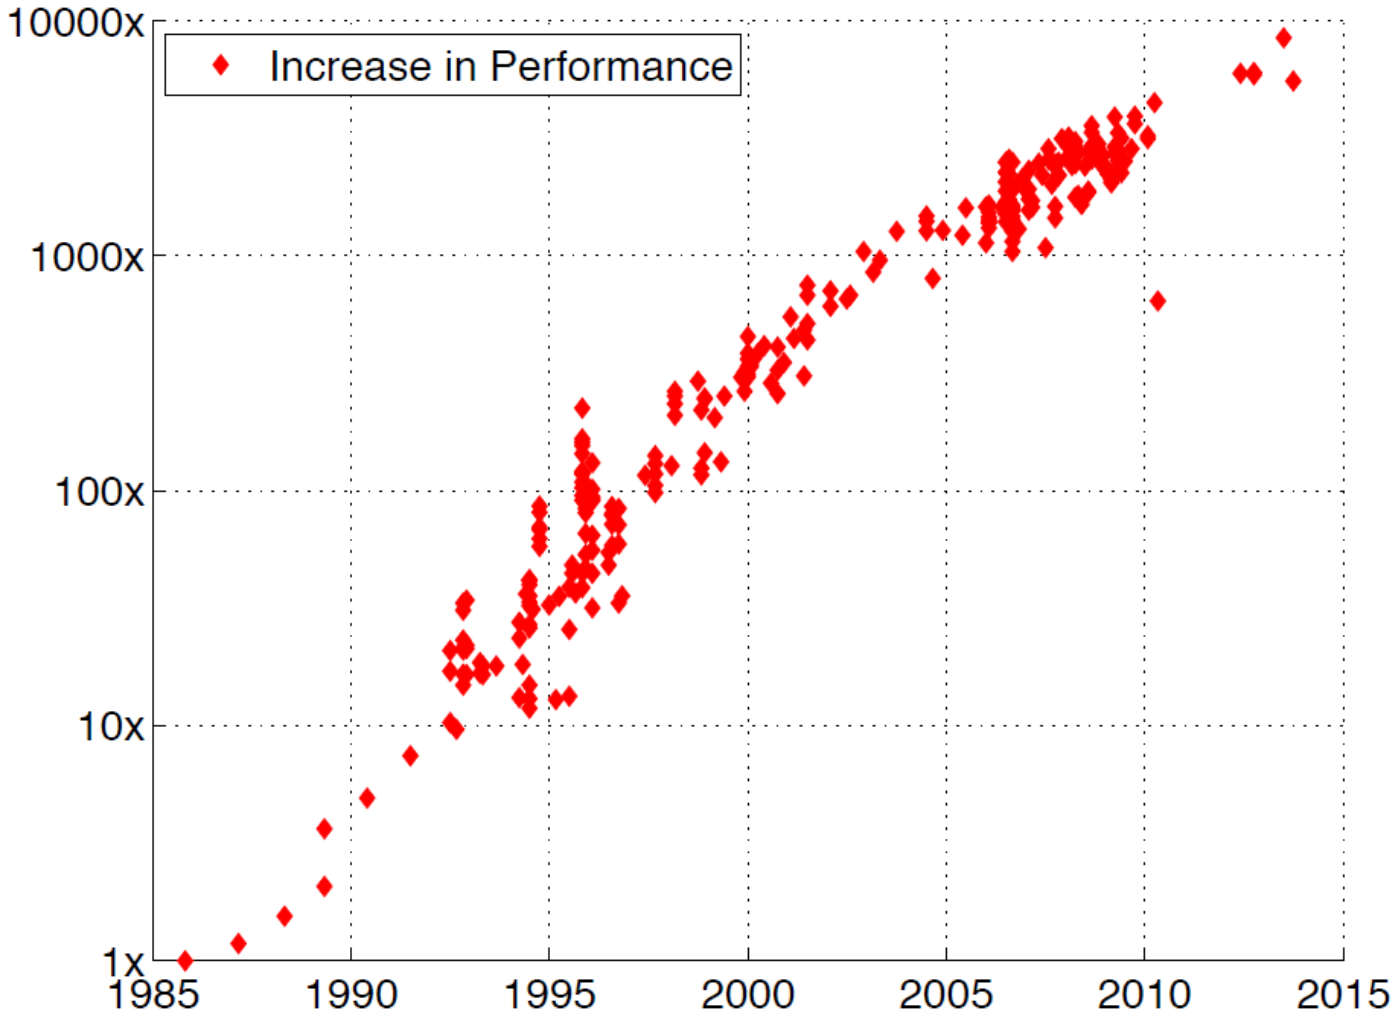

http://cpudb.stanford.edu/

Houston, We Have A Problem

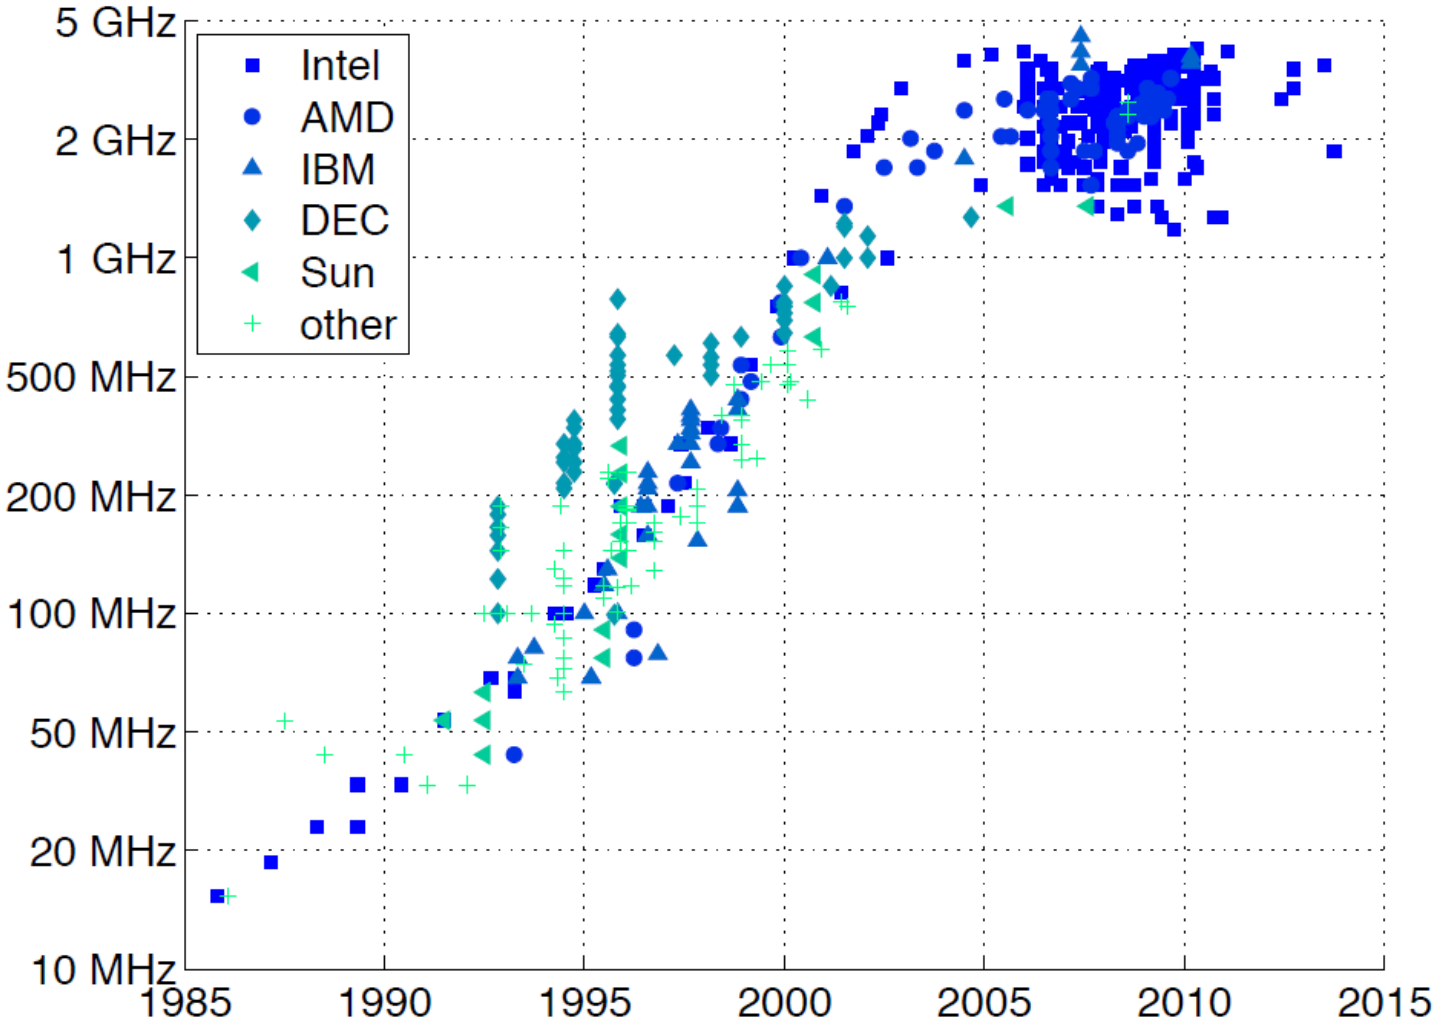

http://cpudb.stanford.edu/

The Power Limit

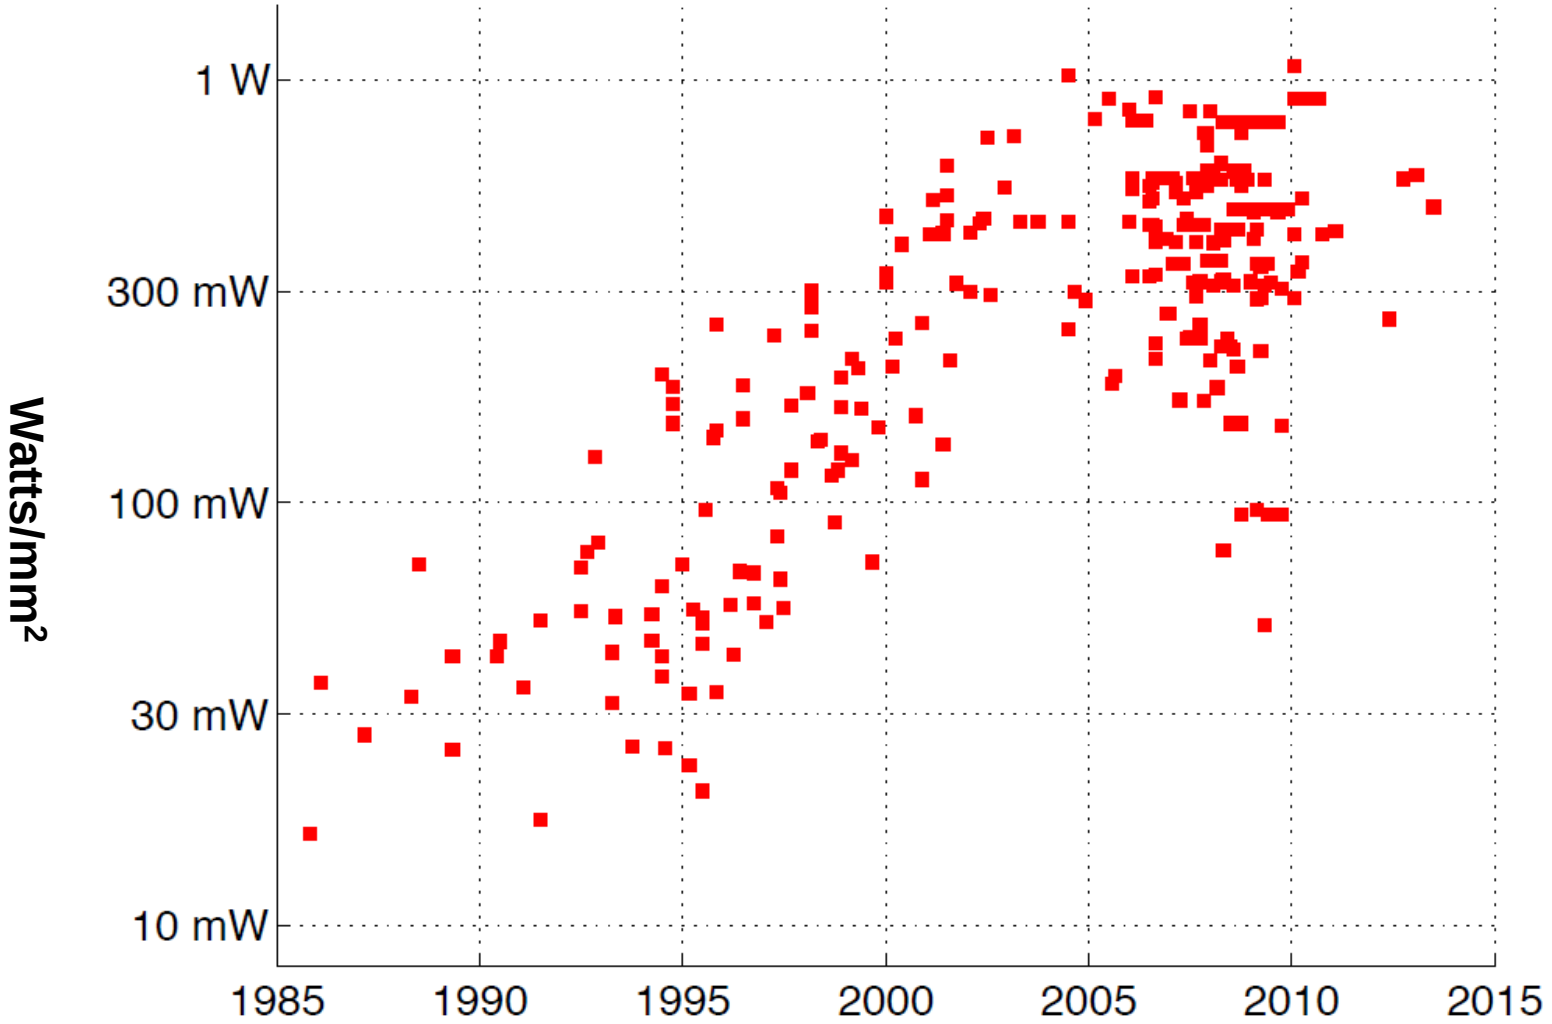

http://cpudb.stanford.edu/

Power Increased Because We Were ~~Greedy~~ Clever

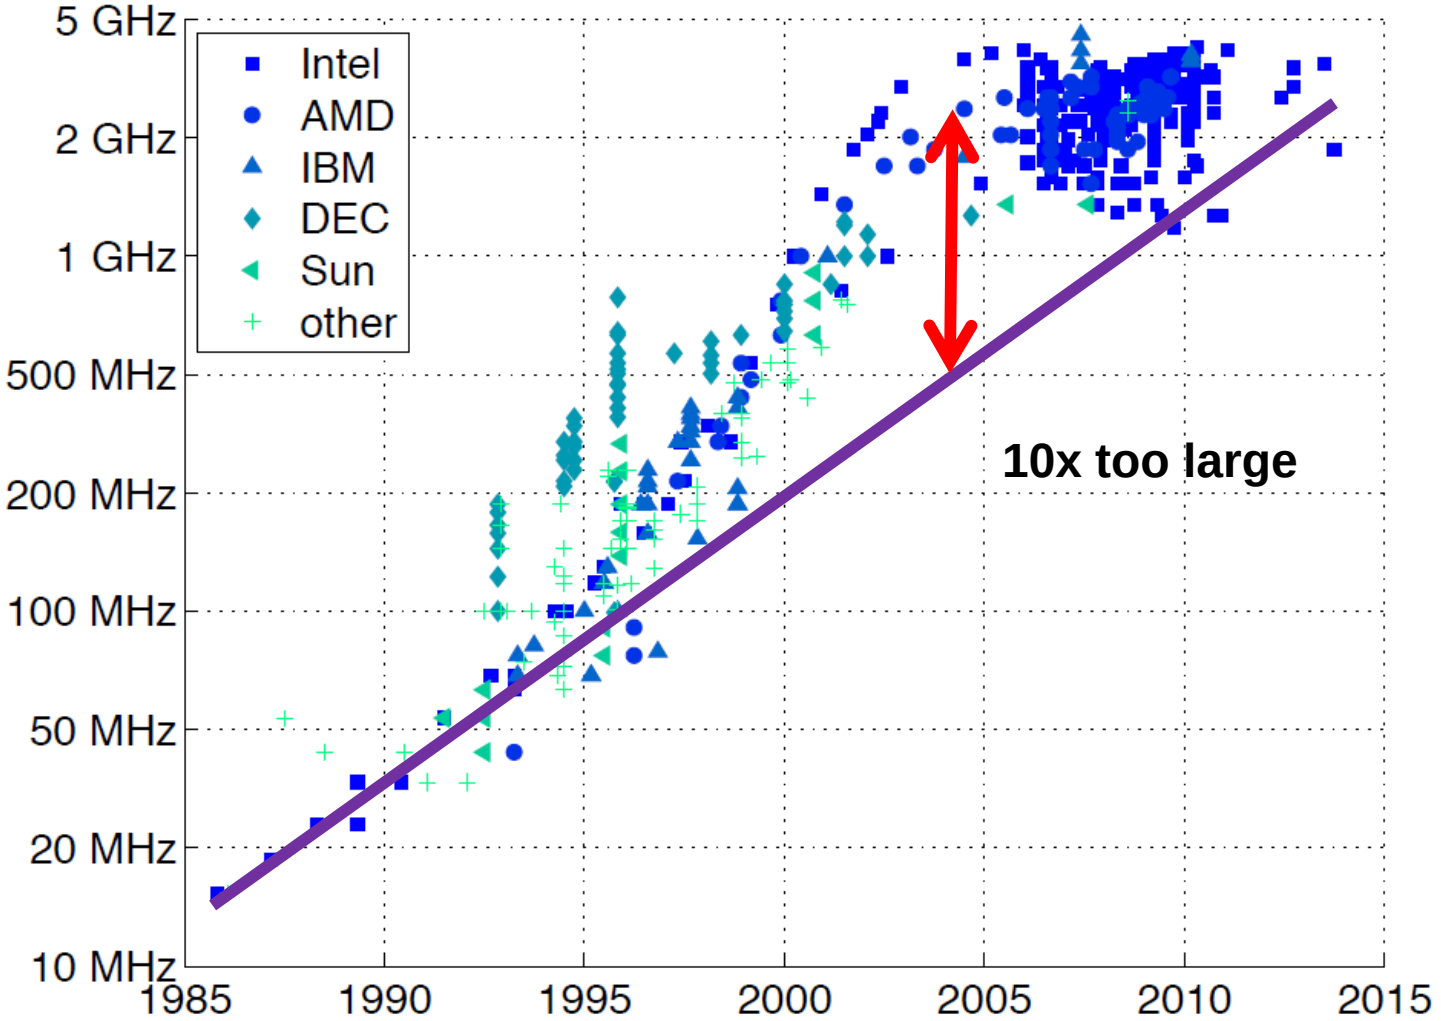

http://cpudb.stanford.edu/

This Power Problem Is Not Going Away: $P = \alpha C * Vdd^2 * f$

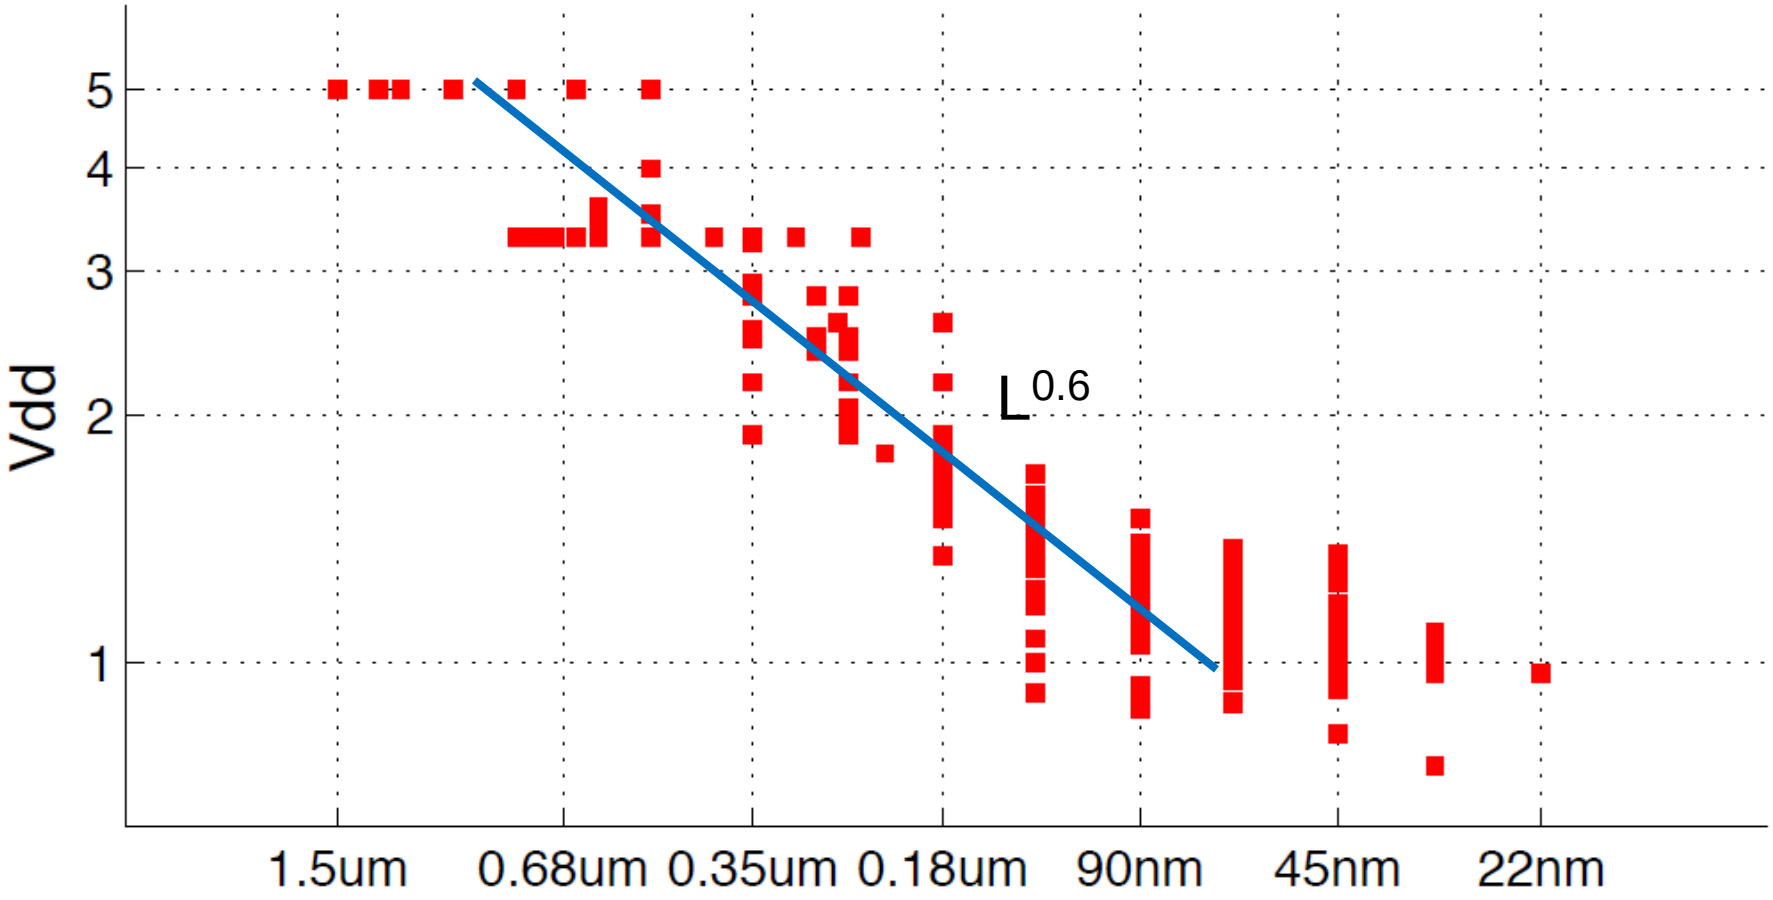

http://cpudb.stanford.edu/

Think About It

$$P = \frac{\text{ENERGY}}{\text{OP}} \quad \frac{\text{OPS}}{\text{S}}$$

Technology to the Rescue?

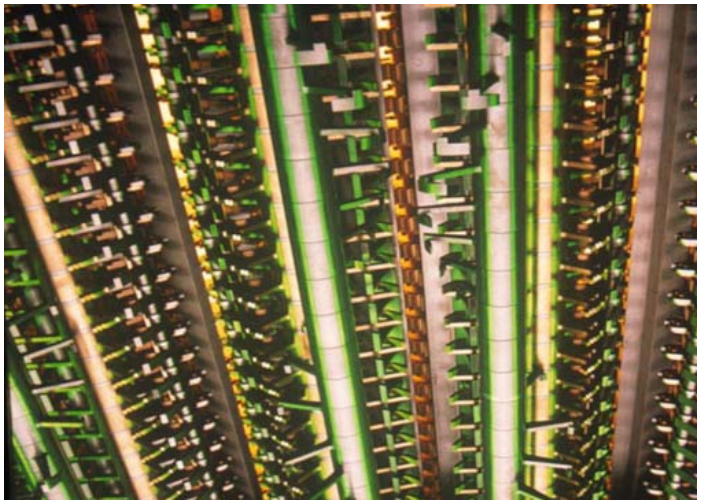

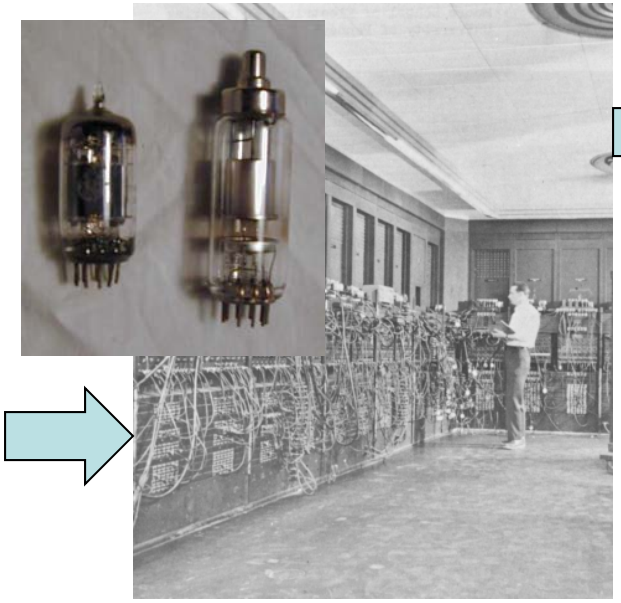

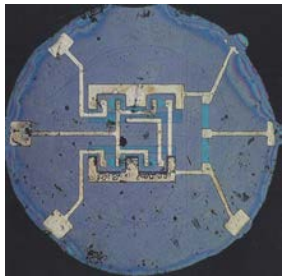

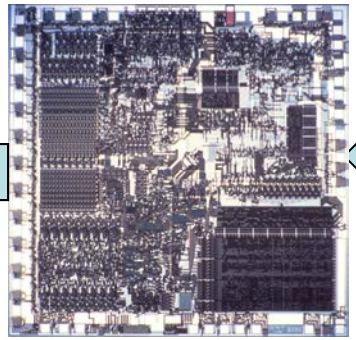

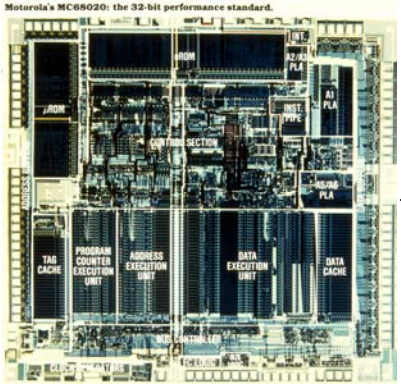

Problems w/ Replacing CMOS

Pretty fundamental physics

- Avoiding this problem will be hard

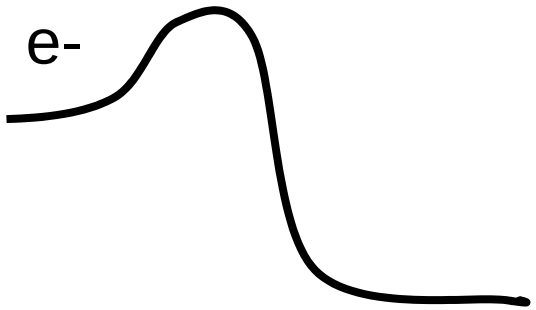

Its capability is pretty amazing

- fJ/gate, 10ps delays, 10^9 working devices

Catch - 22

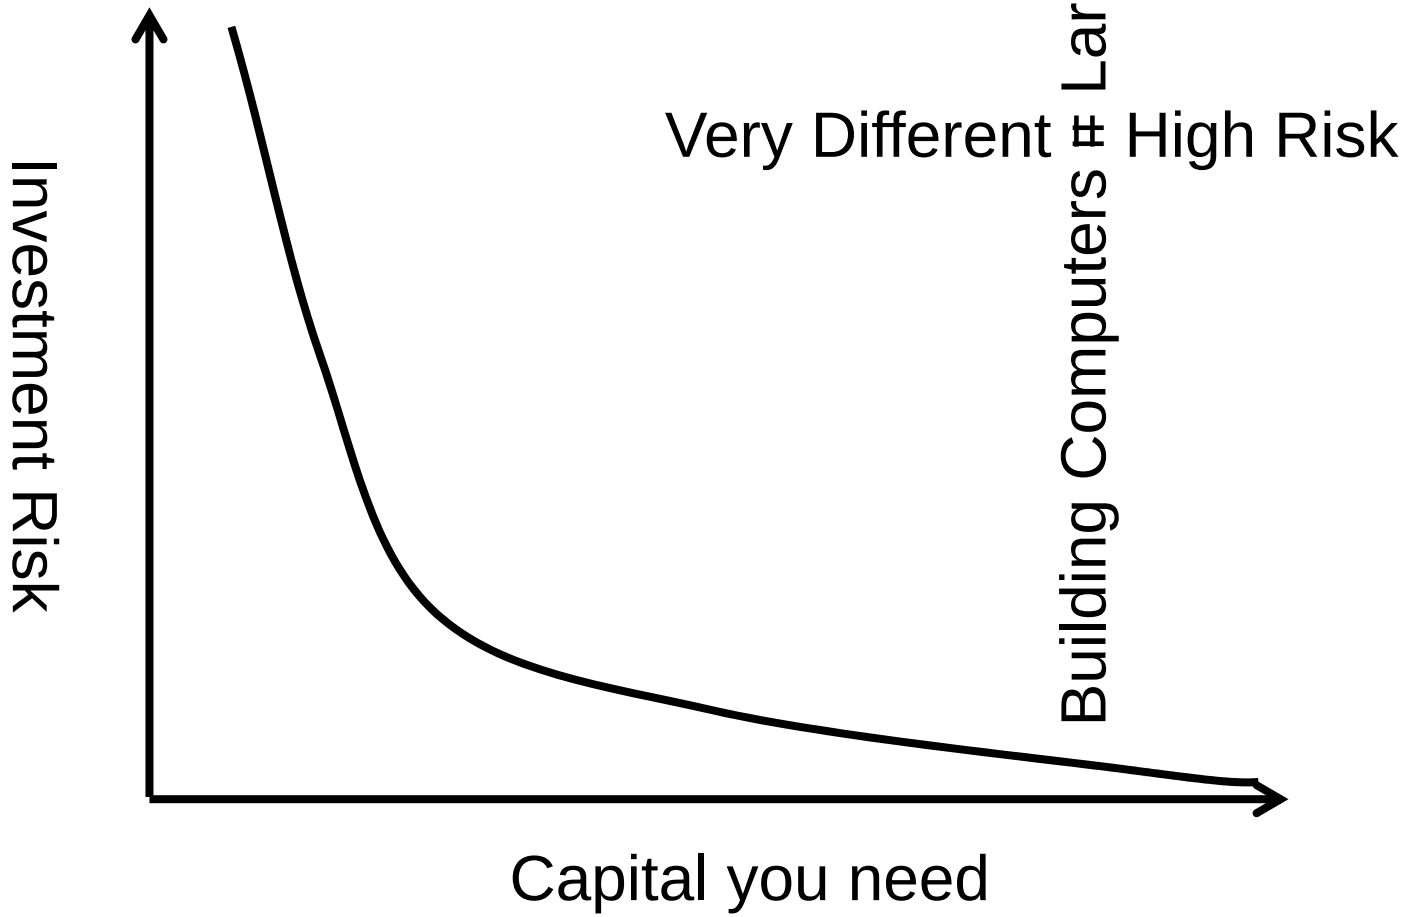

The Truth About Innovation

Start by creating new markets

Our CMOS Future

Will see tremendous innovative uses of computation

- Capability of today's technology is incredible
- Can add computing and communication for nearly $0
- Key questions are what problems need to be solved?

Most performance system will be energy limited

- These systems will be optimized for energy efficiency

Power = Energy/Op * Ops/sec

Processor Energy – Delay Trade-off

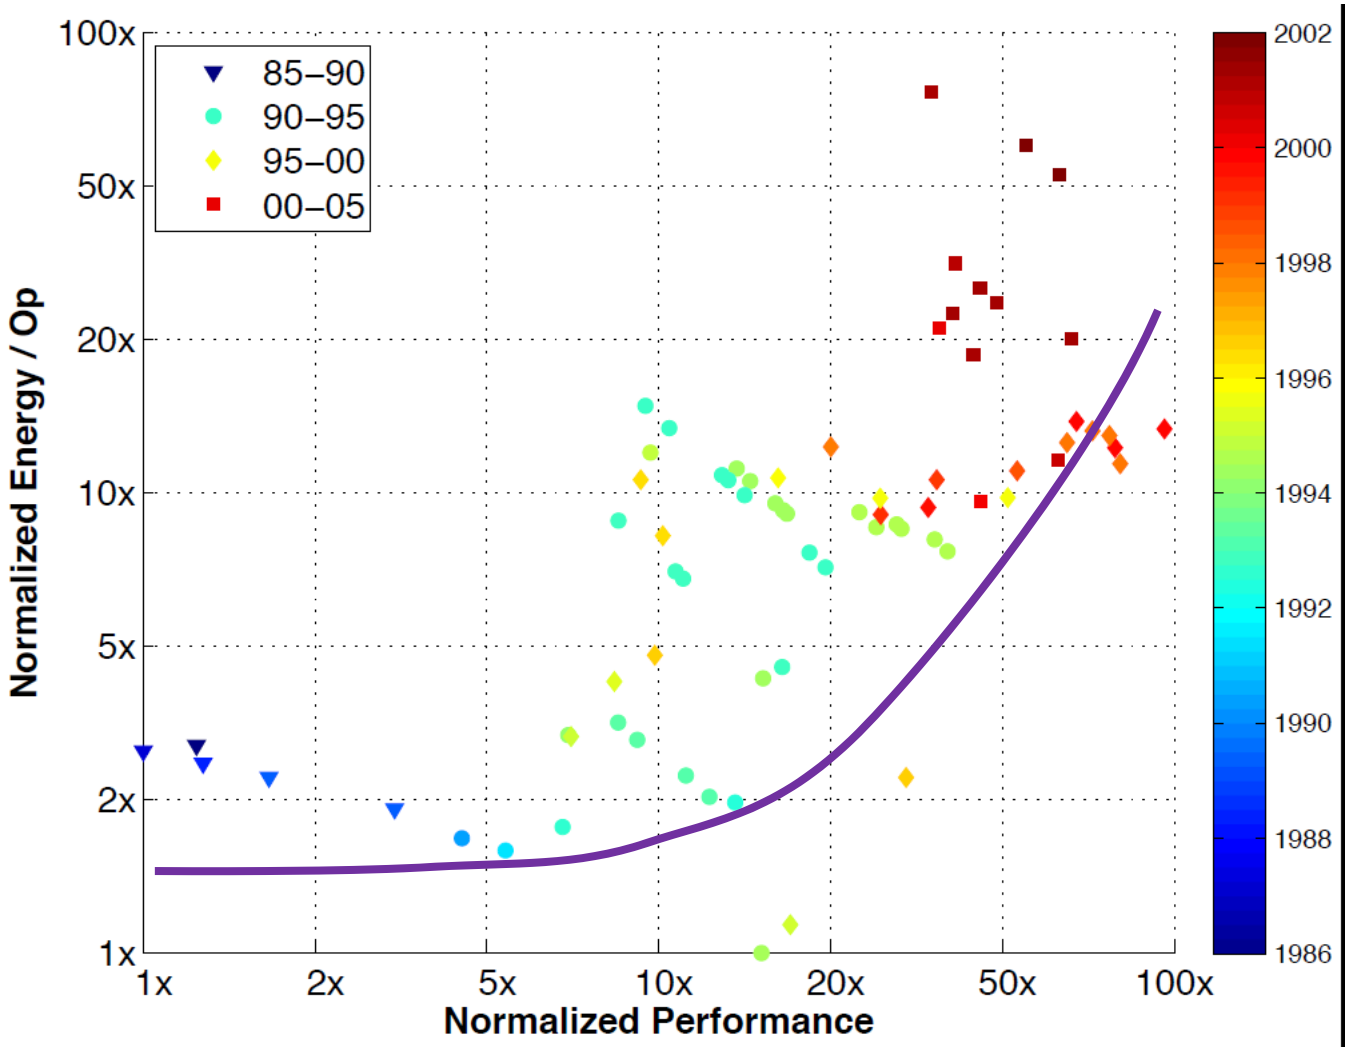

http://cpudb.stanford.edu/

The Rise of Multi-Core Processors

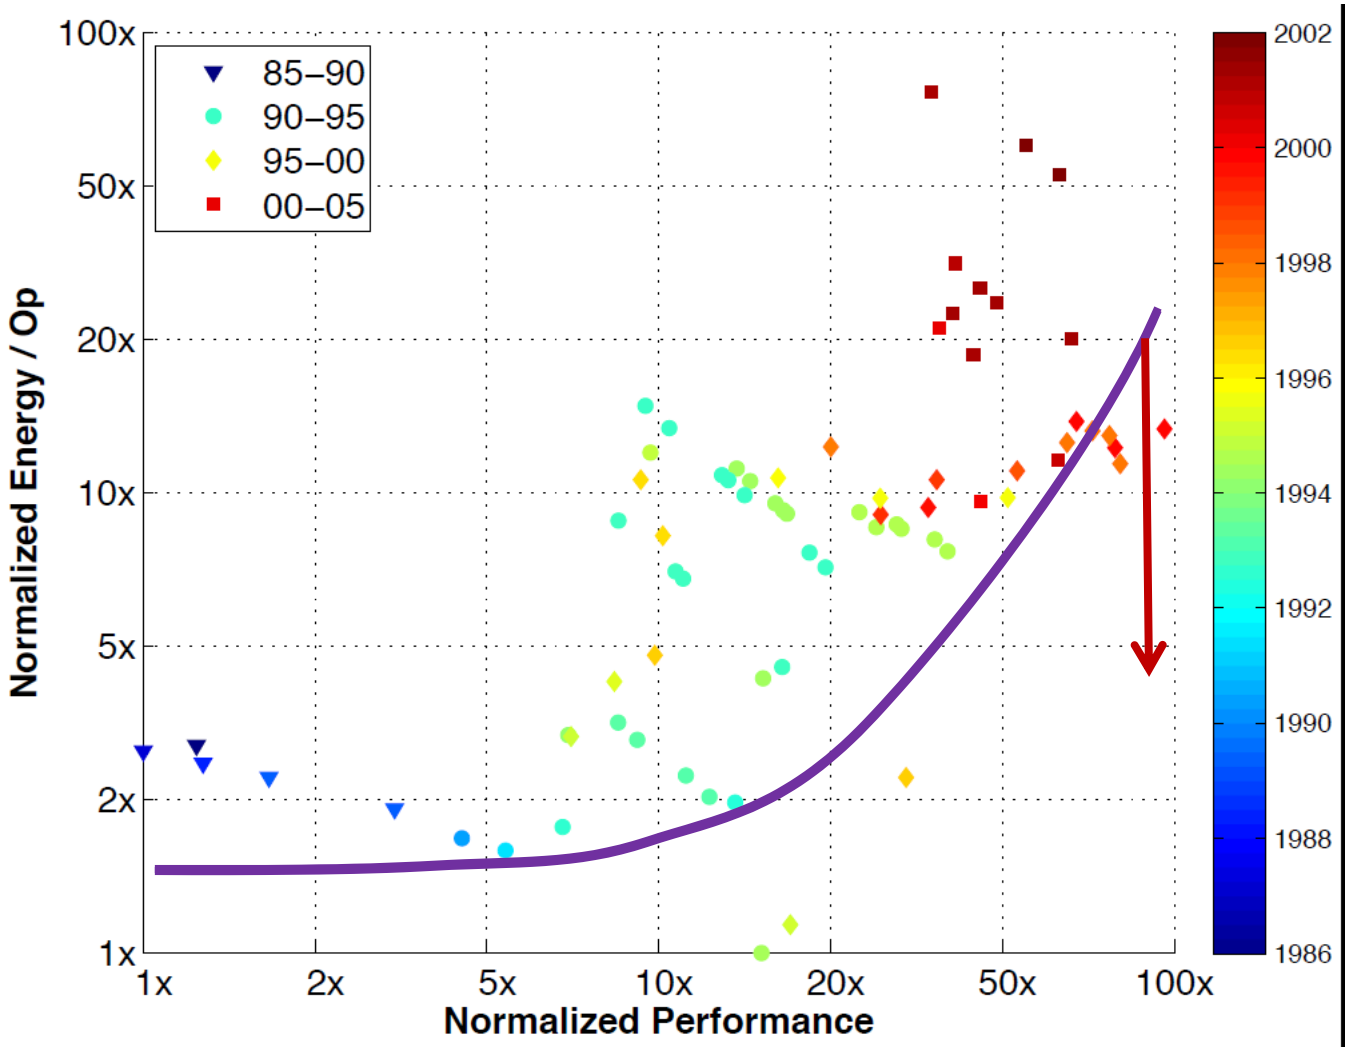

http://cpudb.stanford.edu/

The Stagnation of Multi-Core Processors

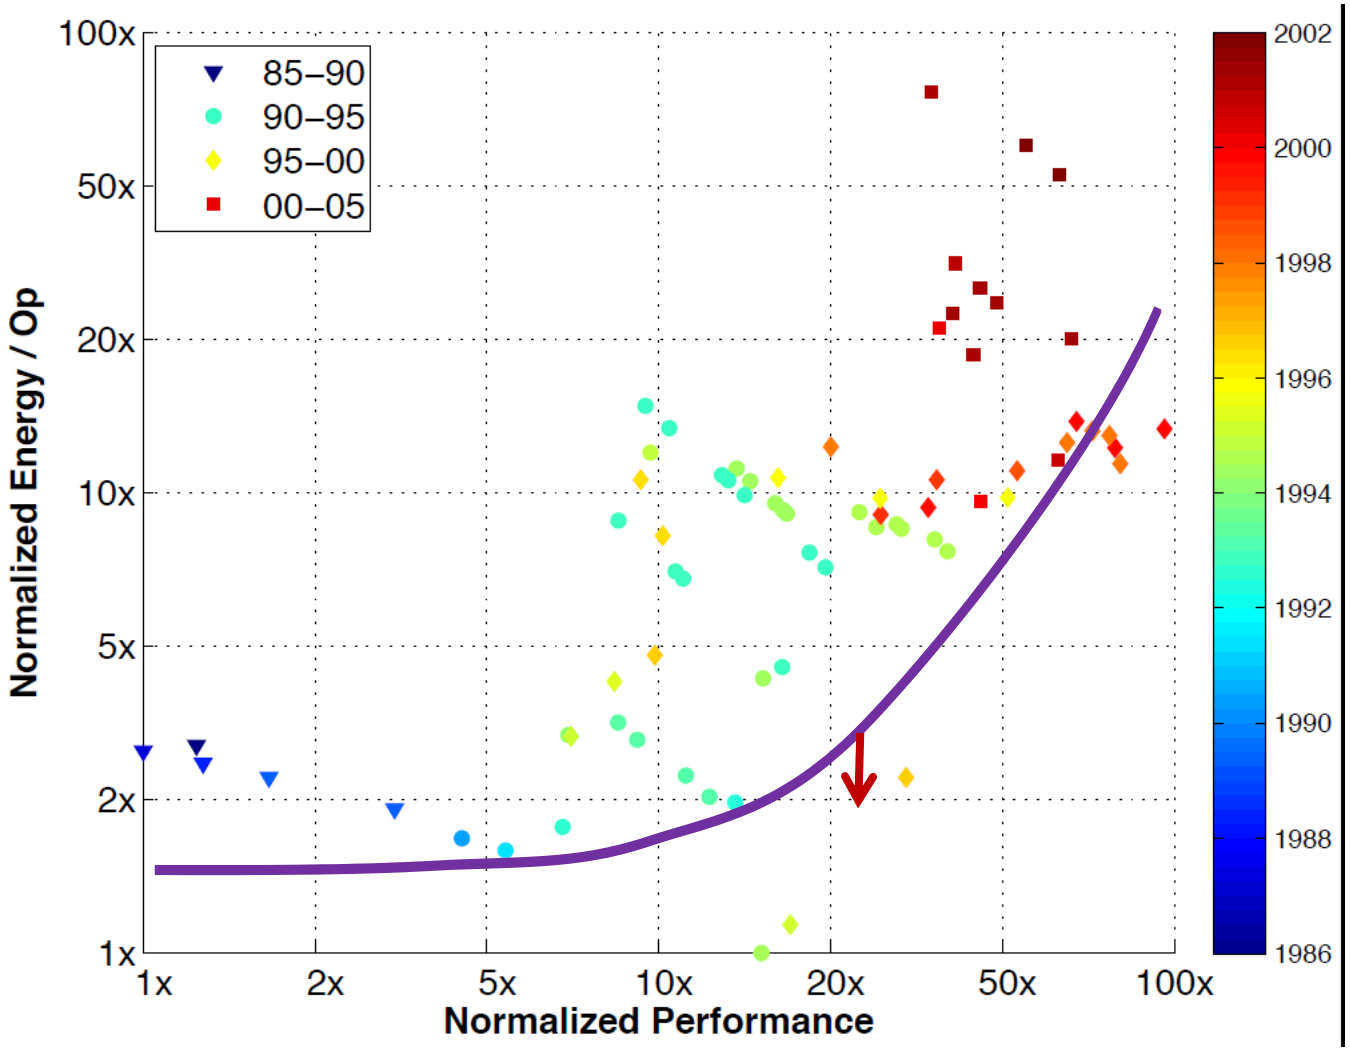

http://cpudb.stanford.edu/

Optimizing Parallel Machines (GPUs)

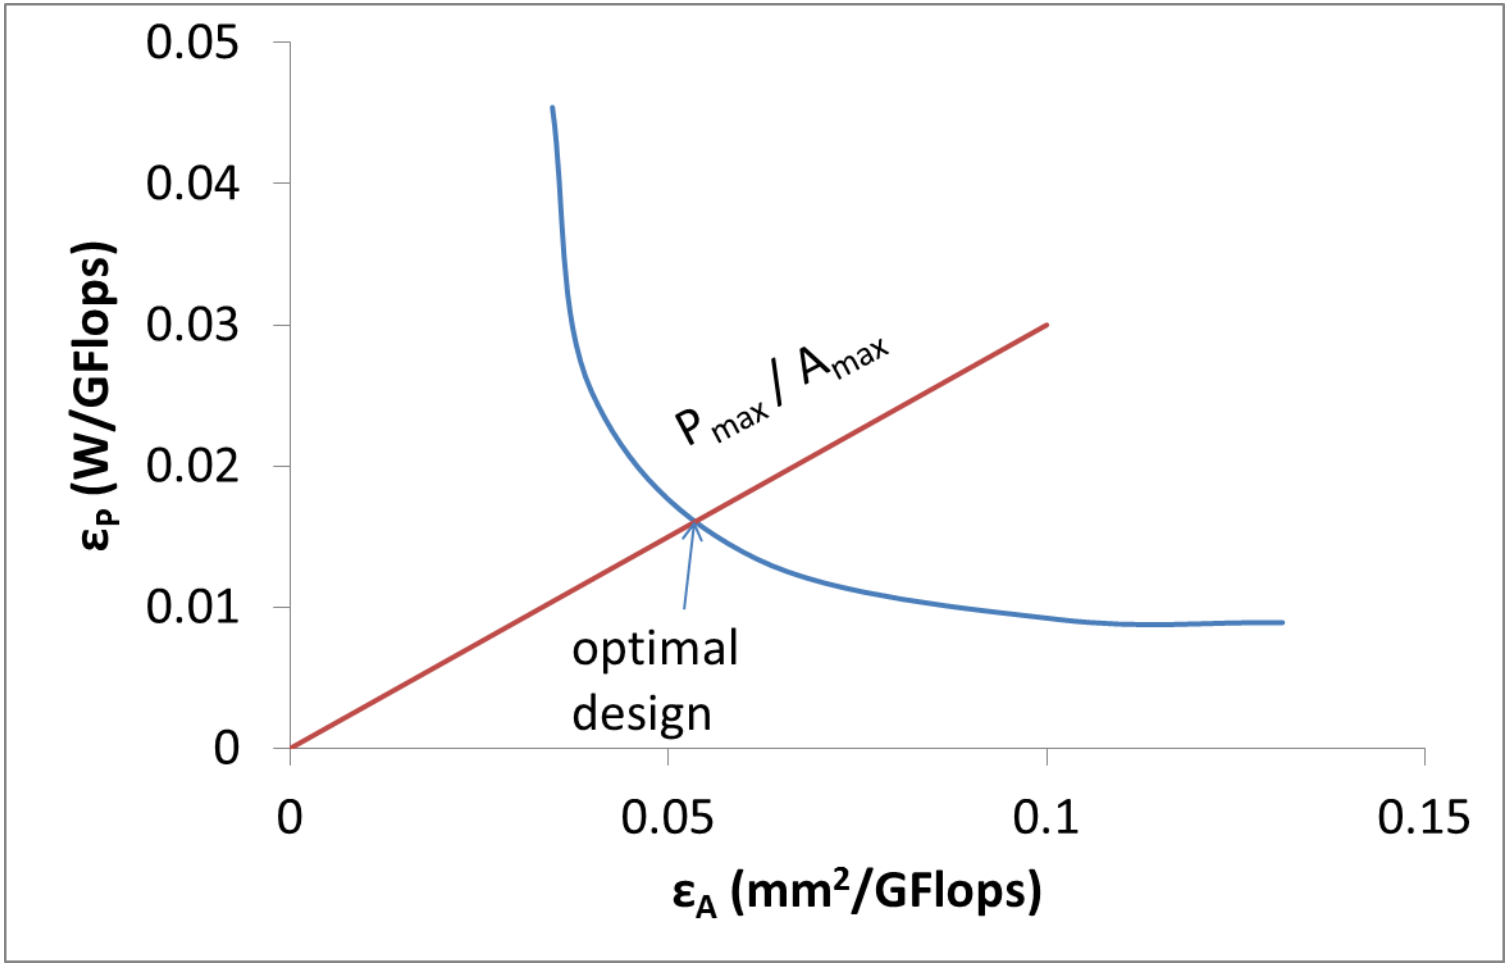

Galal et al, Trans on Computers, July 2011, pp 913,922

Have A Shiny Ball, Now What?

Signal Processing ASICs

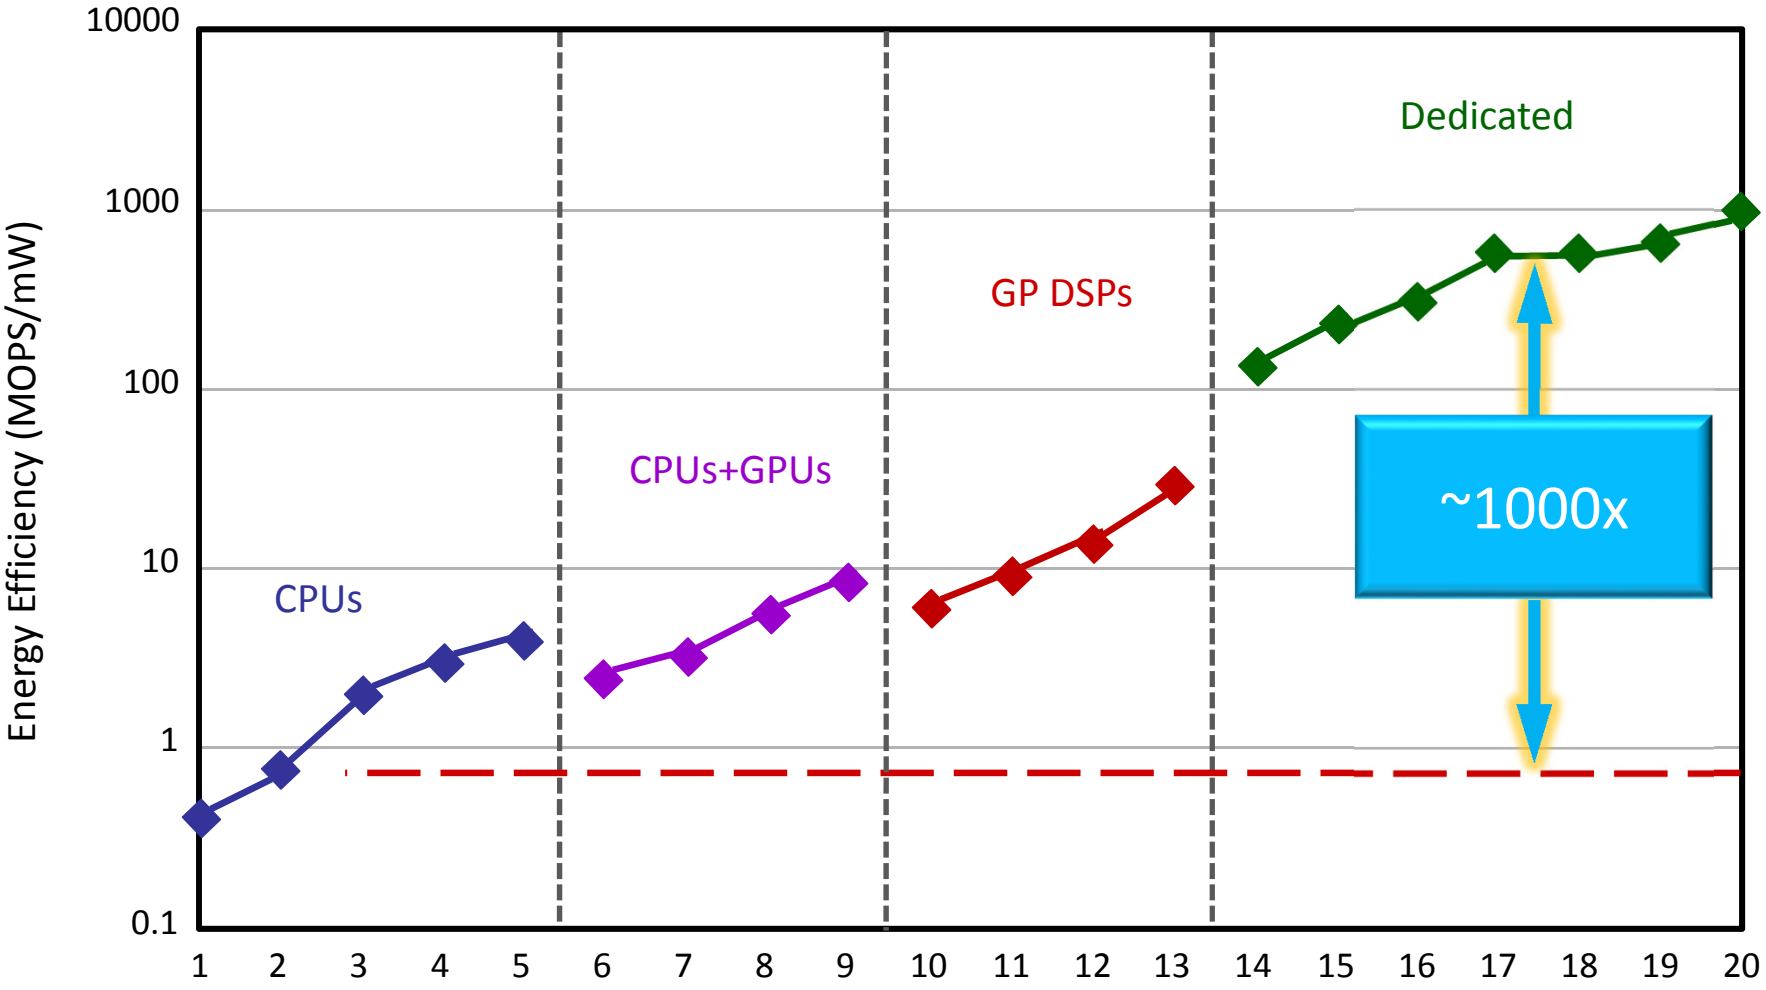

Markovic, EE292 Class, Stanford, 2013

The Push For Specialized Hardware

Dark Silicon and the End of Multi

Hadi Esmaeilzadeh† Emily Blem‡ Renée St. Amant§

†University of Washington ‡The University of Texas at

hadianeh@cs.washington.edu blem@cs.wisc.edu

ABSTRACT

Since 2005, processor designers have increase
ploit Moore's Law scaling, rather than focusing
formance. The failure of Dennard scaling, to which
ticore parts is partially a response, may soon limit mu
just as single-core scaling has been curtailed. This p
multicore scaling limits by combining device scaling,
scaling, and multicore scaling to measure the speedup pot
a set of parallel workloads for the next five technology gener
For device scaling, we use both the ITRS projections and
of more conservative device scaling parameters. To model sing
core scaling, we combine measurements from over 150 processor
to derive Pareto-optimal frontiers for area/performance and pow
er/performance. Finally, to model multicore scaling, we build a de
tailed performance model of upper-bound performance and lower-
bound core power. The multicore designs we study include single-
threaded CPU-like and massively threaded GPU-like multicore chip
organizations with symmetric, asymmetric, dynamic, and composed
topologies. The study shows that regardless of chip organization
and topology, multicore scaling is power limited to a degree not
widely appreciated by the computing community. Even at 22 nm
(just one year from now), 21% of a fixed-size chip must be powered
off, and at 8 nm, this number grows to more than 50%. Through
2024, only 7.9× average speedup is possible across commonly used
parallel workloads, leaving a nearly 24-fold gap from a target of
doubled performance per generation.

Categories and Subject Descriptors: C.0 [Computer Systems Or-
ganization] General — Modeling of computer architecture; C.0
[Computer Systems Organization] General — System architectures

General Terms: Design, Measurement, Performance

Keywords: Dark Silicon, Modeling, Power, Technology Scaling,
Multicore

Conservation Cores:
Reducing the Energy of Mature Computations

Ganesh Venkatesh Jack Sampson Nathan Goulding Saturnino Garcia
Vladyslav Bryksin Jose Lugo-Martinez Steven Swanson Michael Bedford Taylor

Department of Computer Science & Engineering
University of California, San Diego

{gvenkatesh,jsampson,ngouldin,sat,vbryksin,jlugomar,swanson,mbtaylor}@cs.ucsd.edu

Abstract

Growing transistor counts, limited power budgets, and the breakdown of voltage scaling are currently conspiring to create a *utilization wall* that limits the fraction of a chip that can run at full speed at one time. In this regime, specialized, energy-efficient processors can increase parallelism by reducing the per-computation power requirements and allowing more computations to execute under the same power budget. To pursue this goal, this paper introduces *conservation cores*. Conservation cores, or *c-cores*, are specialized processors that focus on reducing energy and energy-delay instead of increasing performance. This focus on energy makes c-cores an excellent match for many applications that would be poor candidates for hardware acceleration (e.g., irregular integer codes). We present a toolchain for automatically synthesizing c-cores from application source code and demonstrate that they can significantly reduce energy and energy-delay for a wide range of applications. The c-cores support patching, a form of targeted reconfigurability, that allows a move to new versions of the software they target. Our results show that conservation cores can reduce energy consumption by 11× and by up to 2.1× for whole applications and the useful lifetime of individual processors.

power. Consequently, the rate at which we can switch transistors is far outpacing our ability to dissipate the heat created by those transistors.

The result is a technology-imposed *utilization wall* that limits the fraction of the chip we can use at full speed at one time. Our experiments with a 45 nm TSMC process show that we can switch less than 7% of a 300mm² die at full frequency within an 80W power budget. ITRS roadmap projections and CMOS scaling theory suggests that this percentage will decrease to less than 3.5% in 32 nm, and will continue to decrease by almost half with each process generation—and even further with 3-D integration.

The effects of the utilization wall are already indirectly apparent in modern processors. Intel's Nehalem provides a "turbo mode" that powers off some cores in order to run others at higher speeds. Another strong indication is that even though native transistor switching speeds have continued to double every two process generations, processor frequencies have not increased substantially over the last 5 years.

In this regime, reducing per-operation energy [19] translates directly into increased potential parallelism for the system: If given computation can be made to consume less power at the same level of performance, other computations can be run in parallel without violating the power budget.

This paper attacks the utilization wall with *conservation cores*. Conservation cores, or *c-cores*, are application-specific hardware circuits created for the purpose of reducing energy consumption in computationally-intensive applications. Since it is no longer possible to run the entire chip at full frequency at once, it makes sense to optimize the portions of the chip that are running so that it is possible for the application at hand. In effect, architects to trade area for energy in order to mitigate the utilization wall has made this trade-off possible. Continued increases in transistor

Before Talking About Specialization

WE SHOULD CHECK ONE MORE THING FIRST

Don’t Forget Memory System Energy

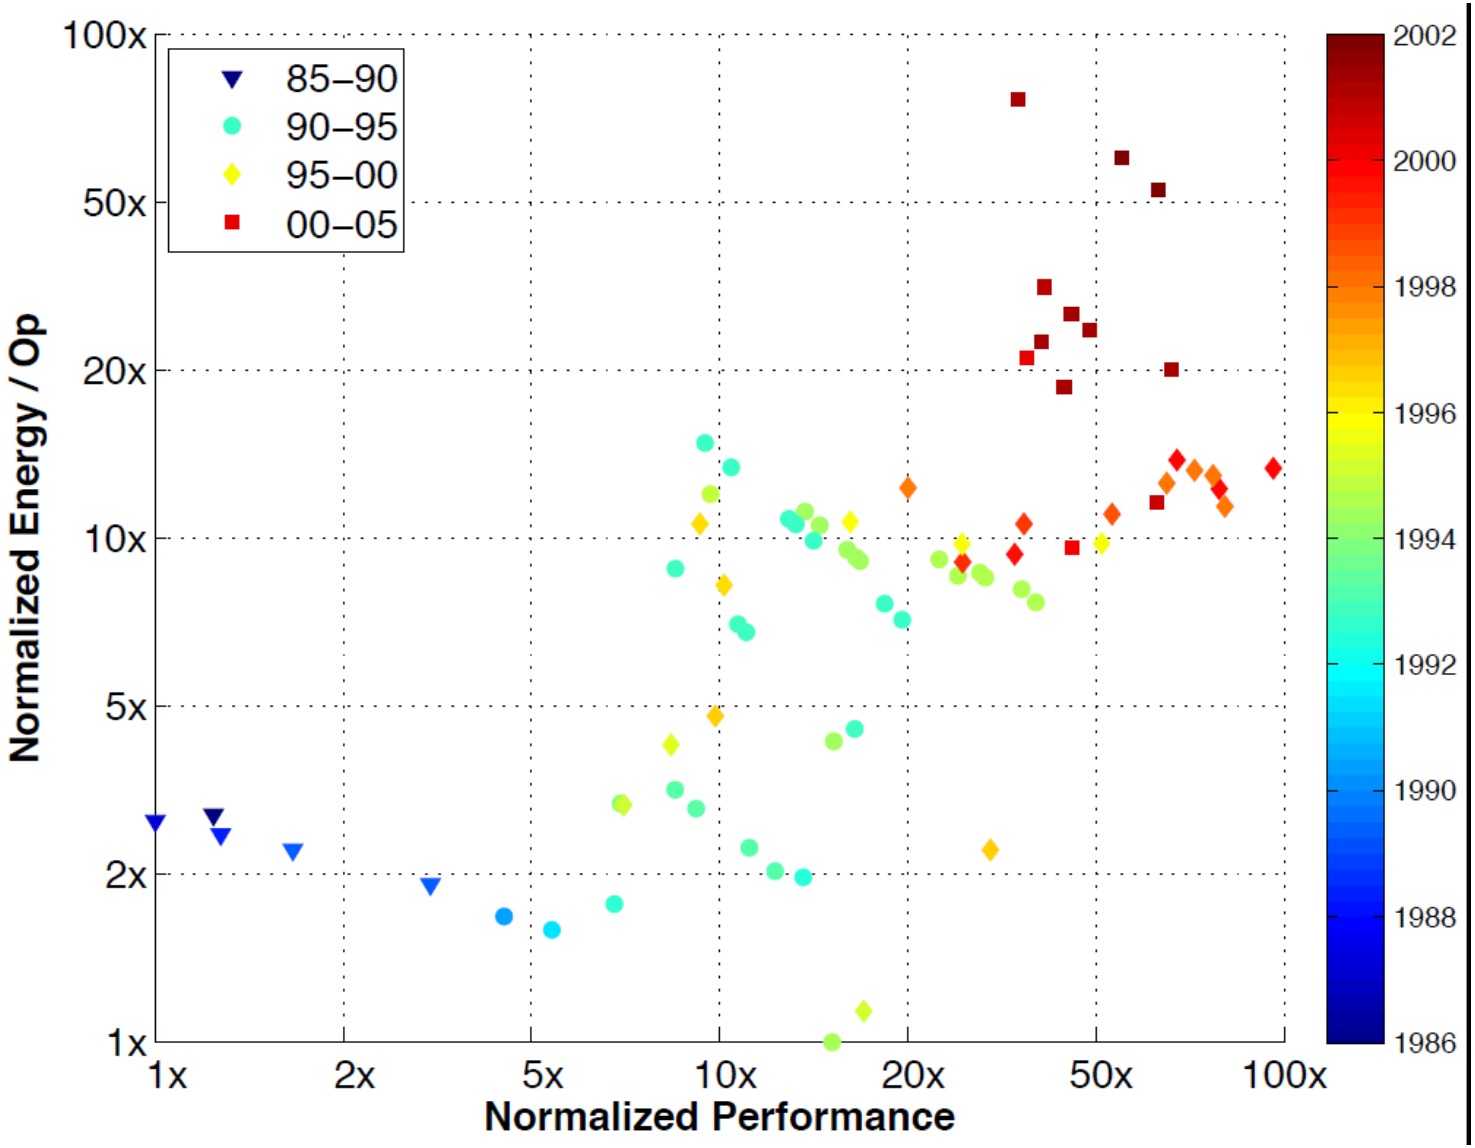

Processor Energy w/ Corrected Cache Sizes

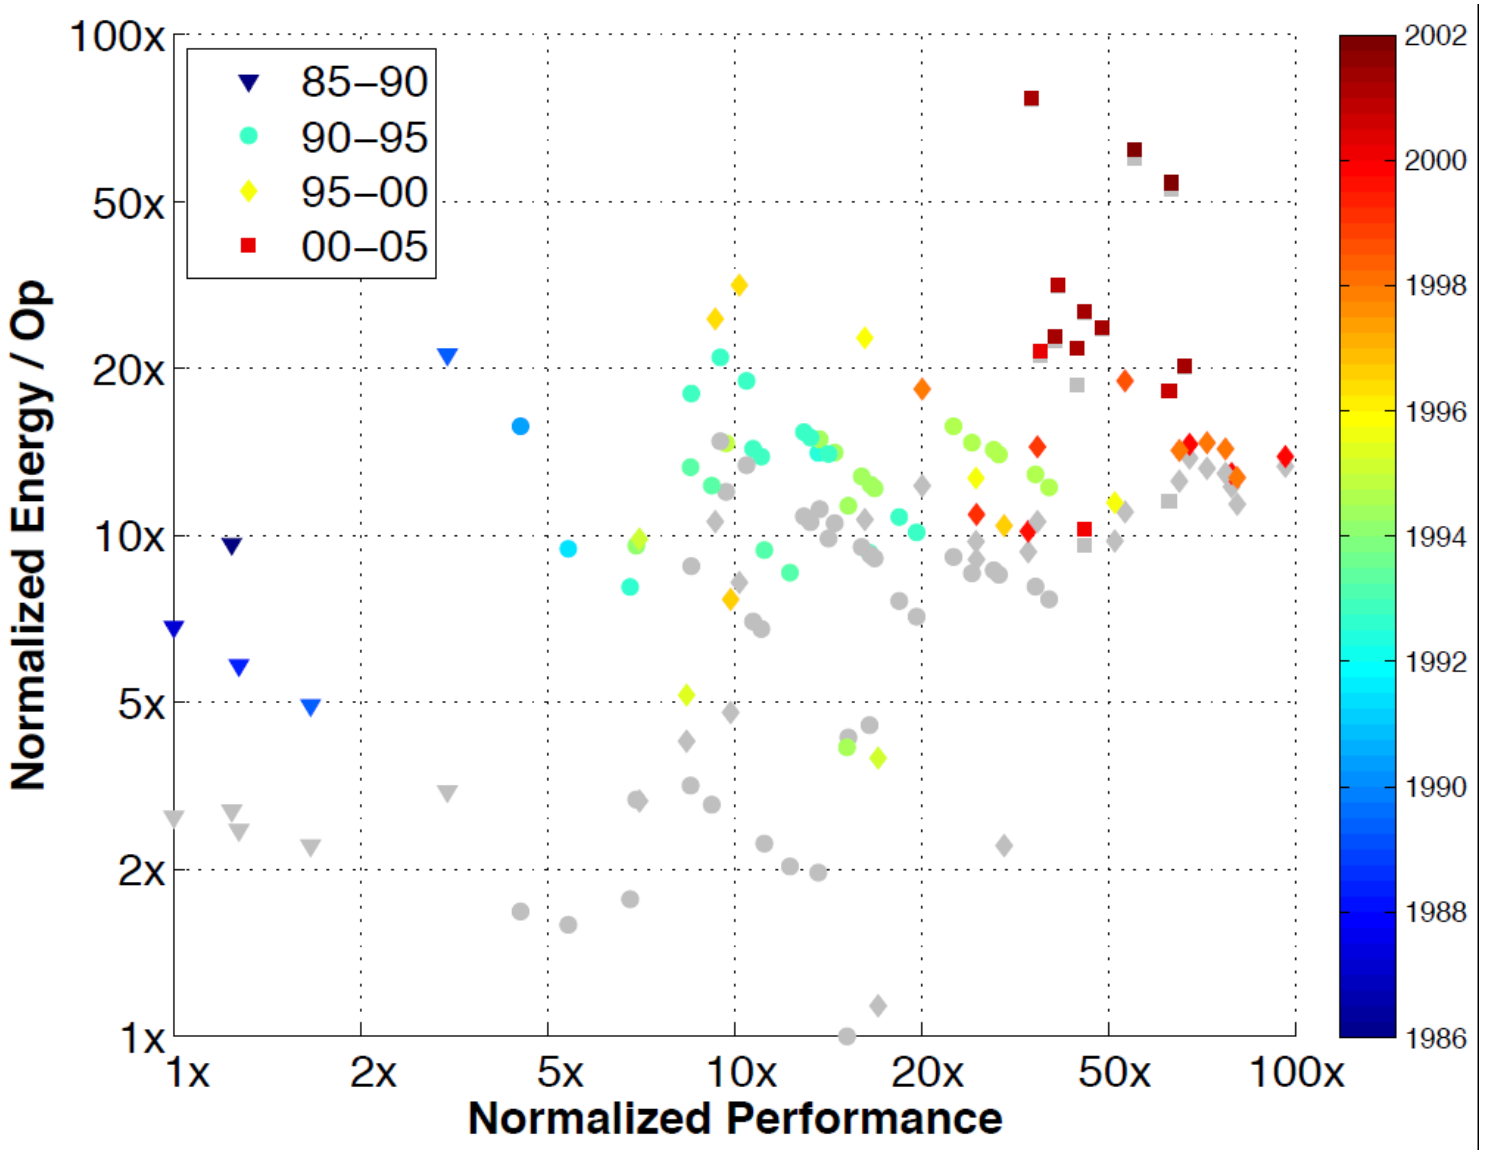

Processor Energy Breakdown

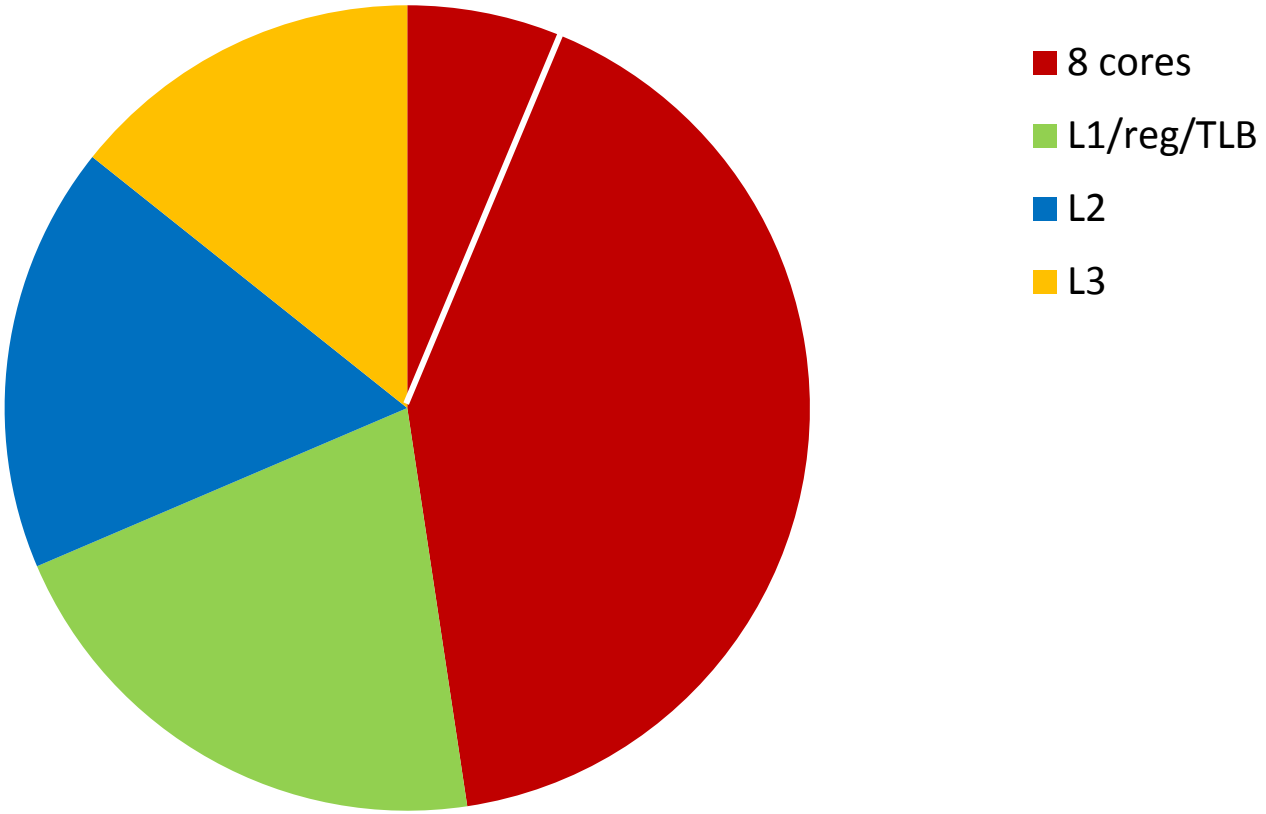

Data Center Energy Specs

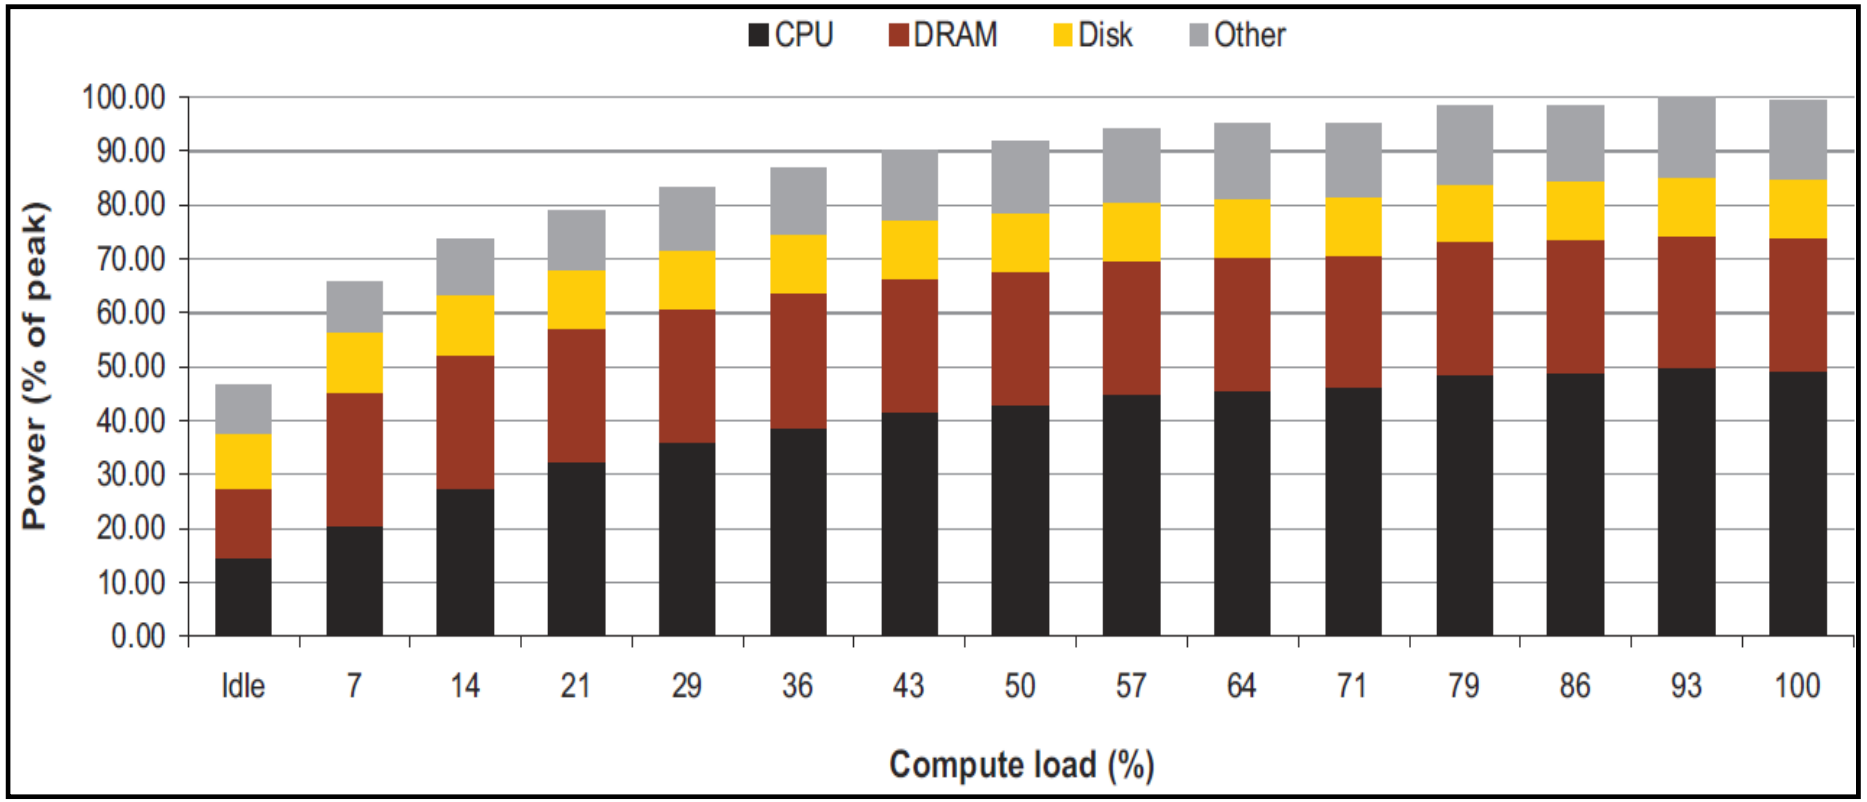

Malladi, ISCA, 2012

SO HOW WILL ACCELERATORS HELP?

What Is Going On Here?

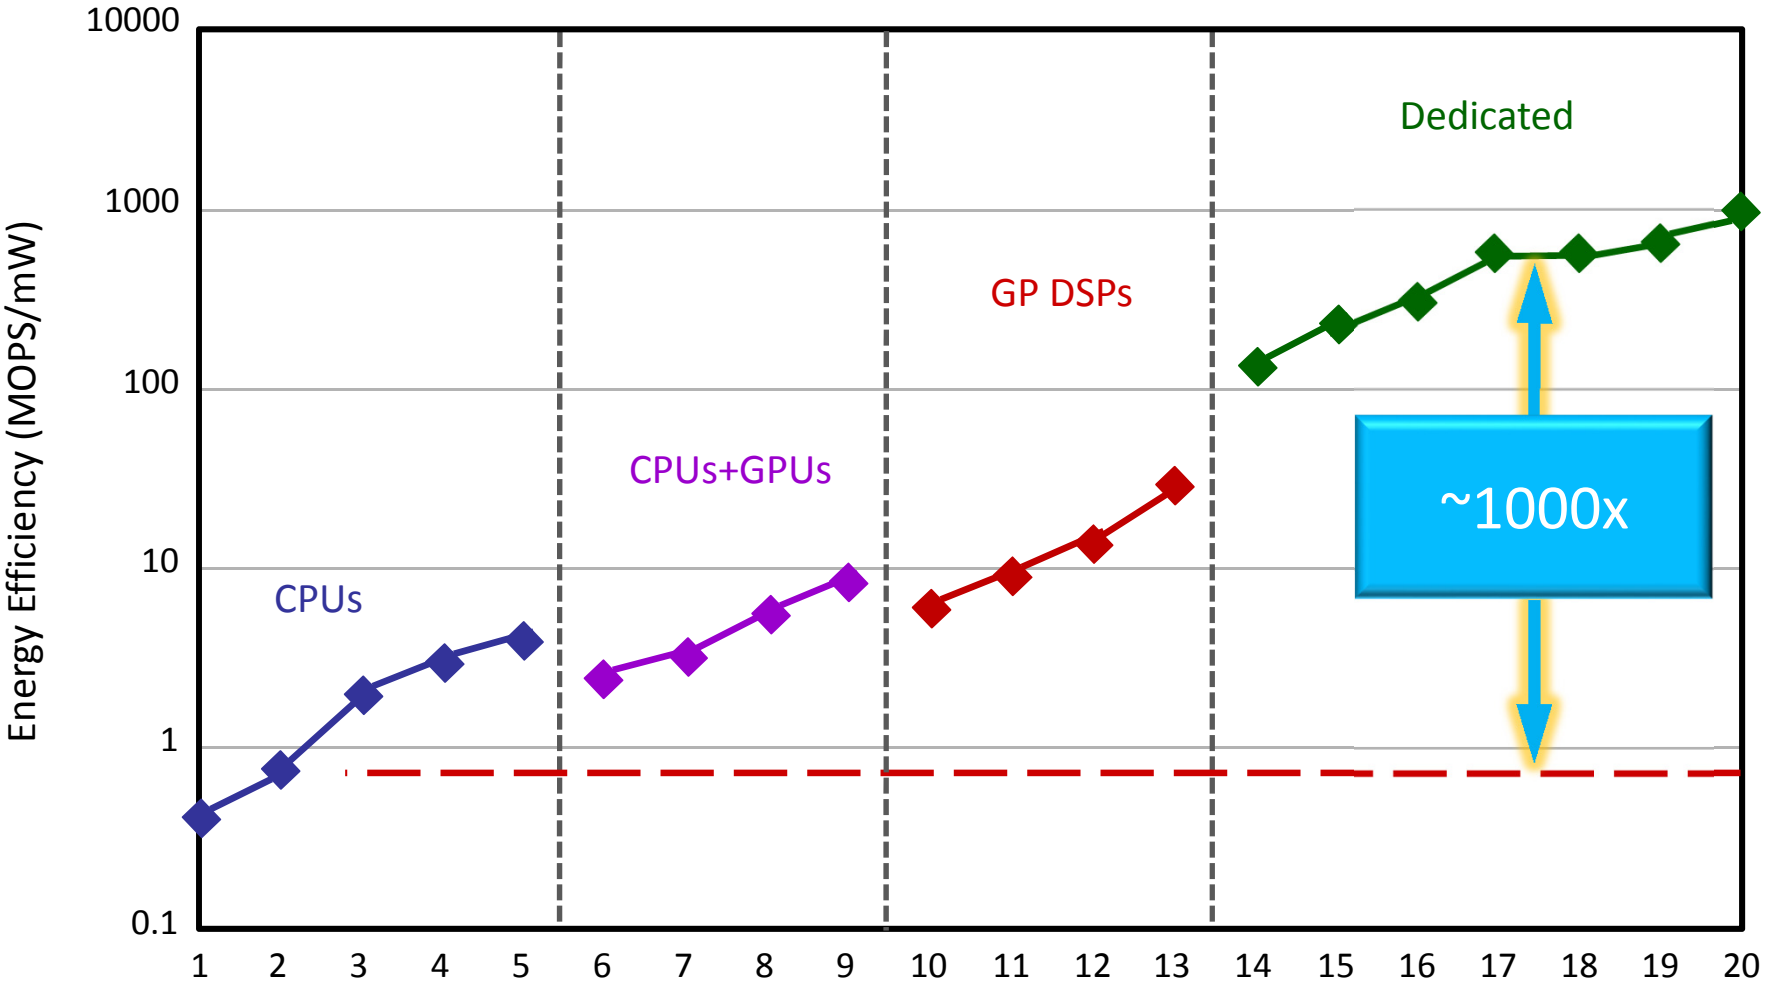

ASIC's Dirty Little Secret

All the ASIC applications have absurd locality

- And work on short integer data

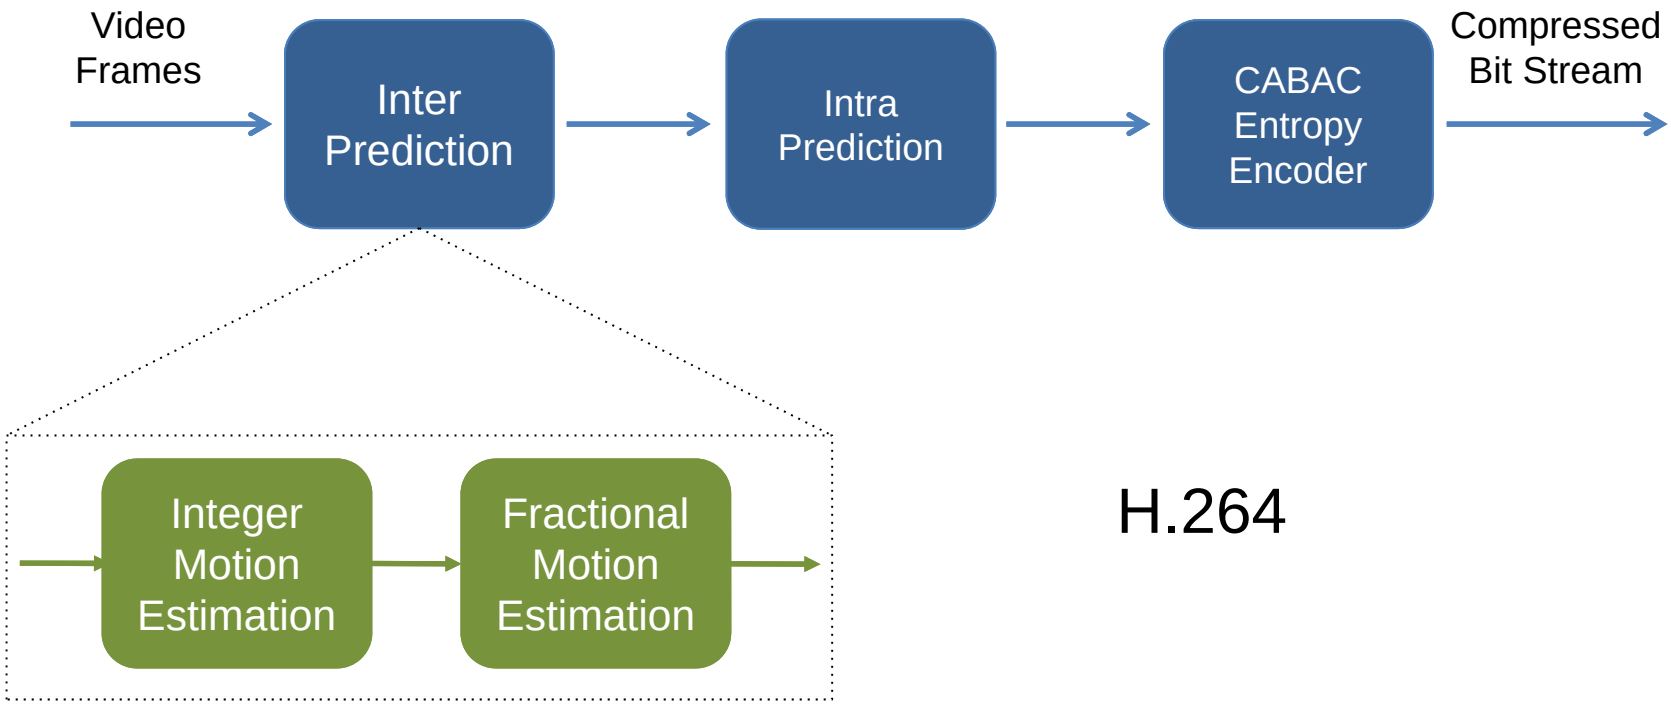

H.264

90% of Execution time is here

Hamid et al, ISCA, 2010

Rough Energy Numbers (45nm)

Integer	
Add	
8 bit	0.03pJ
32 bit	0.1pJ
Mult	
8 bit	0.2pJ
32 bit	3 pJ

FP	
FAdd	
16 bit	0.4pJ
32 bit	0.9pJ
FMult	
16 bit	1pJ
32 bit	4pJ

Memory	
Cache	(64bit)
8KB	10pJ
32KB	20pJ
1MB	100pJ
DRAM	1.3-2.6nJ

Instruction Energy Breakdown

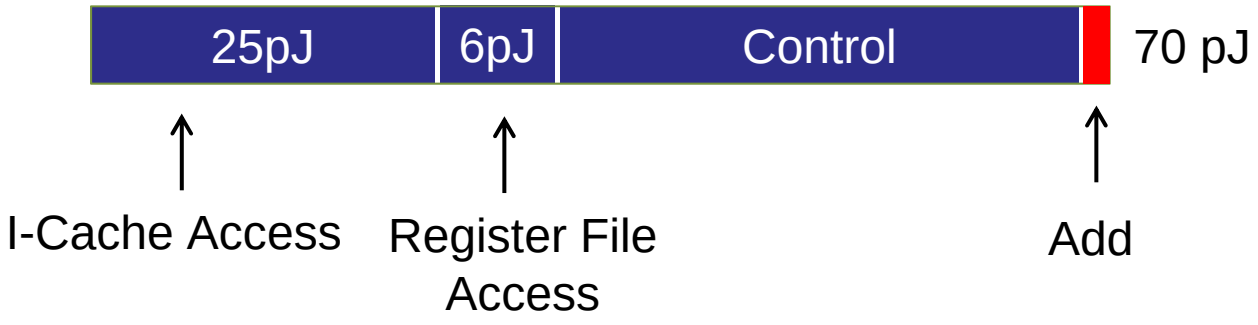

The Truth:
It's More About the Algorithm then the Hardware

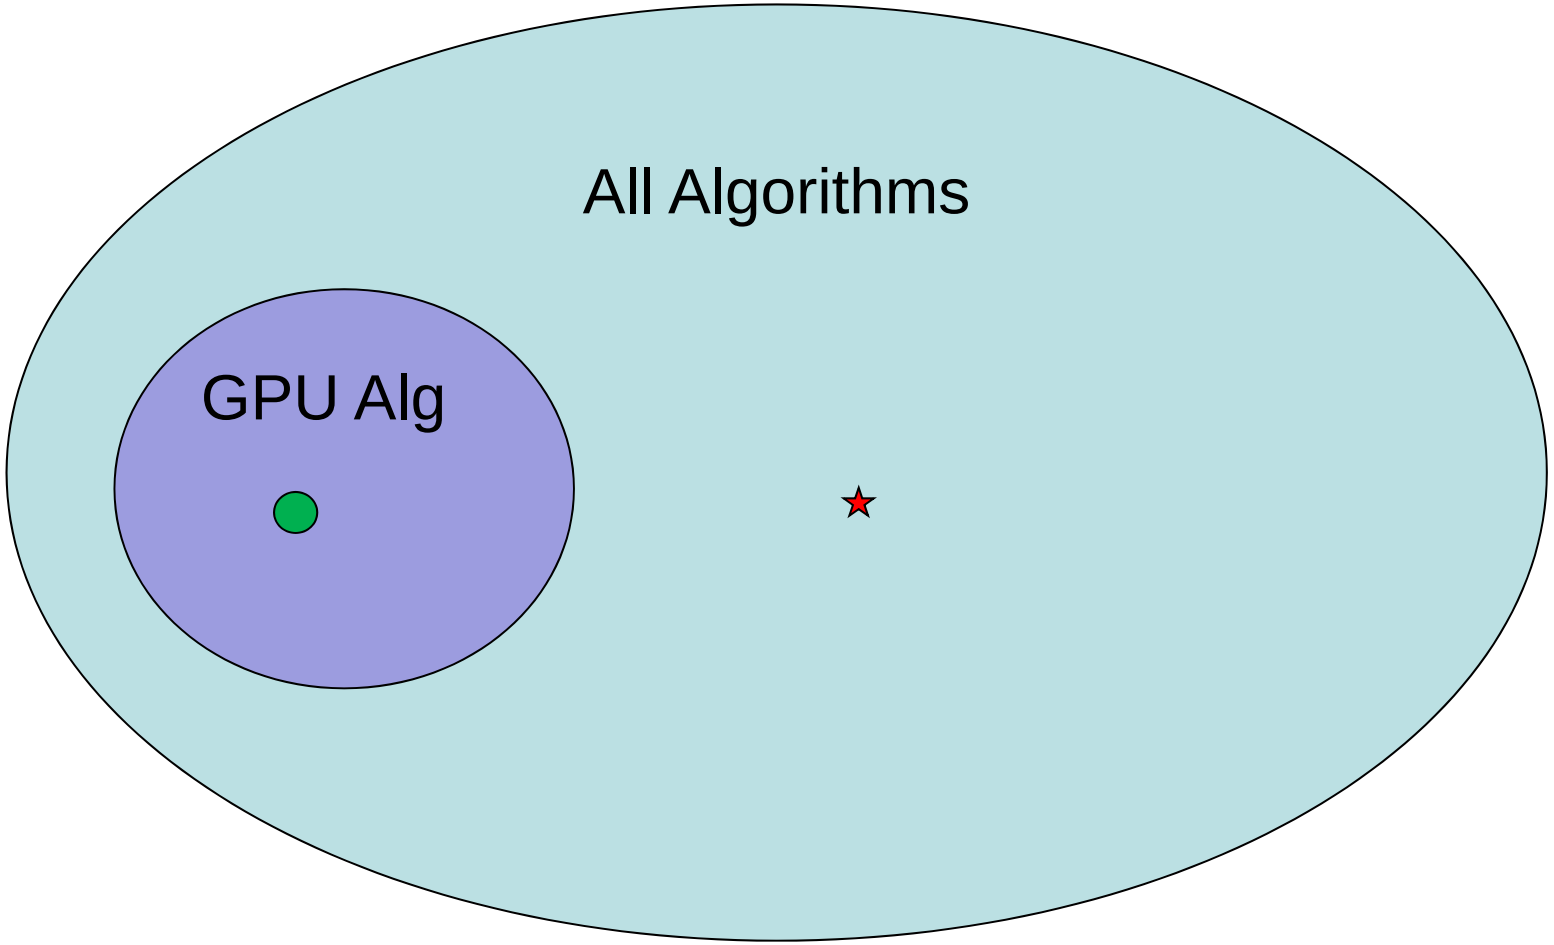

Highly Local Computation Model

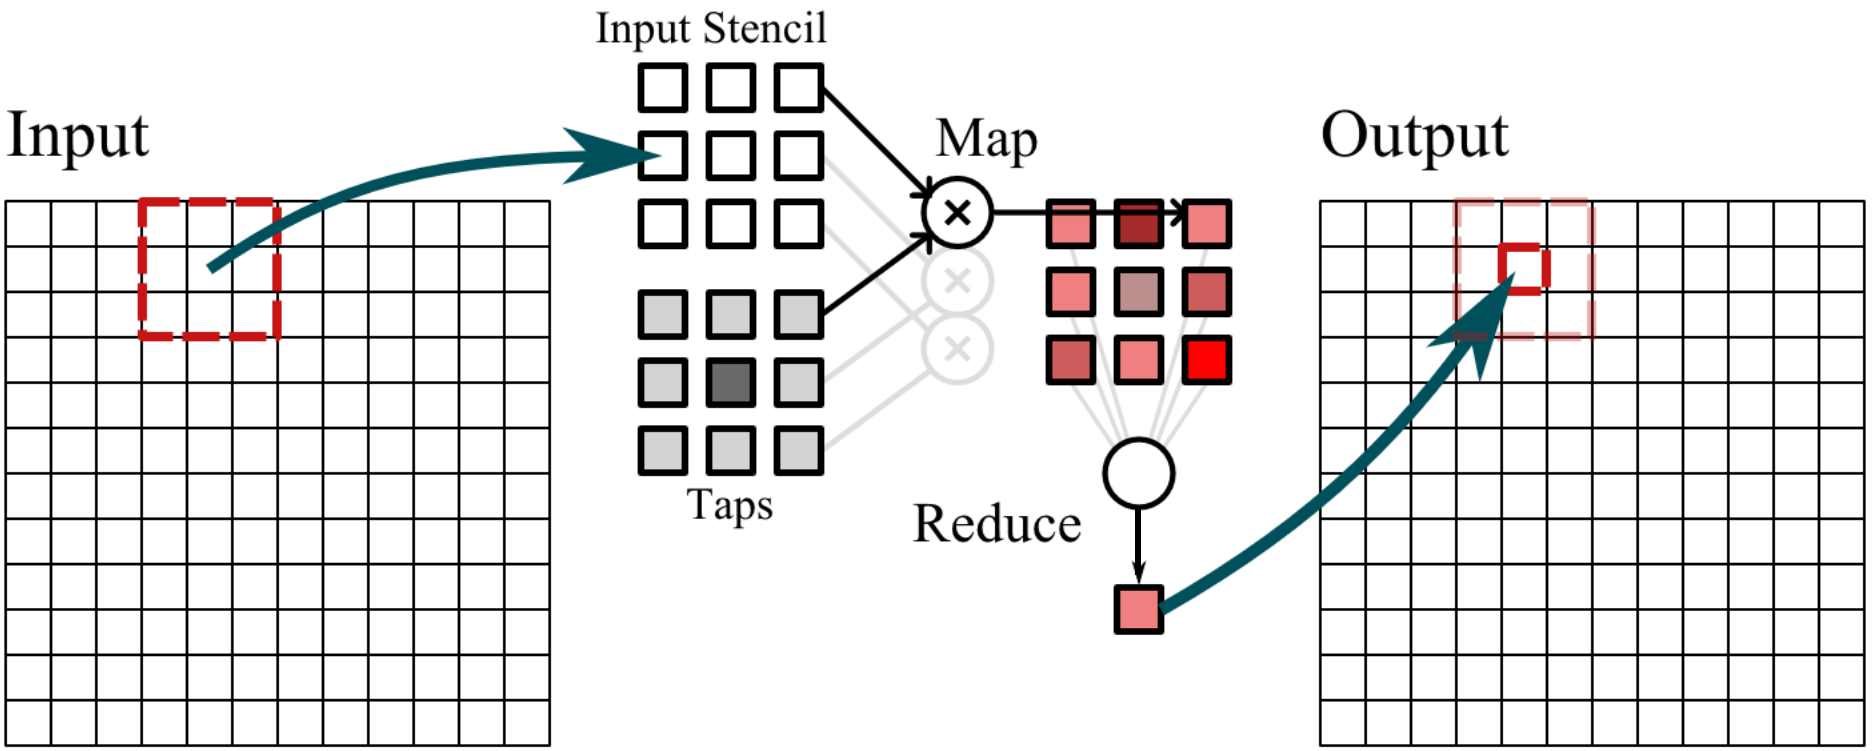

Highly Local Computation Model

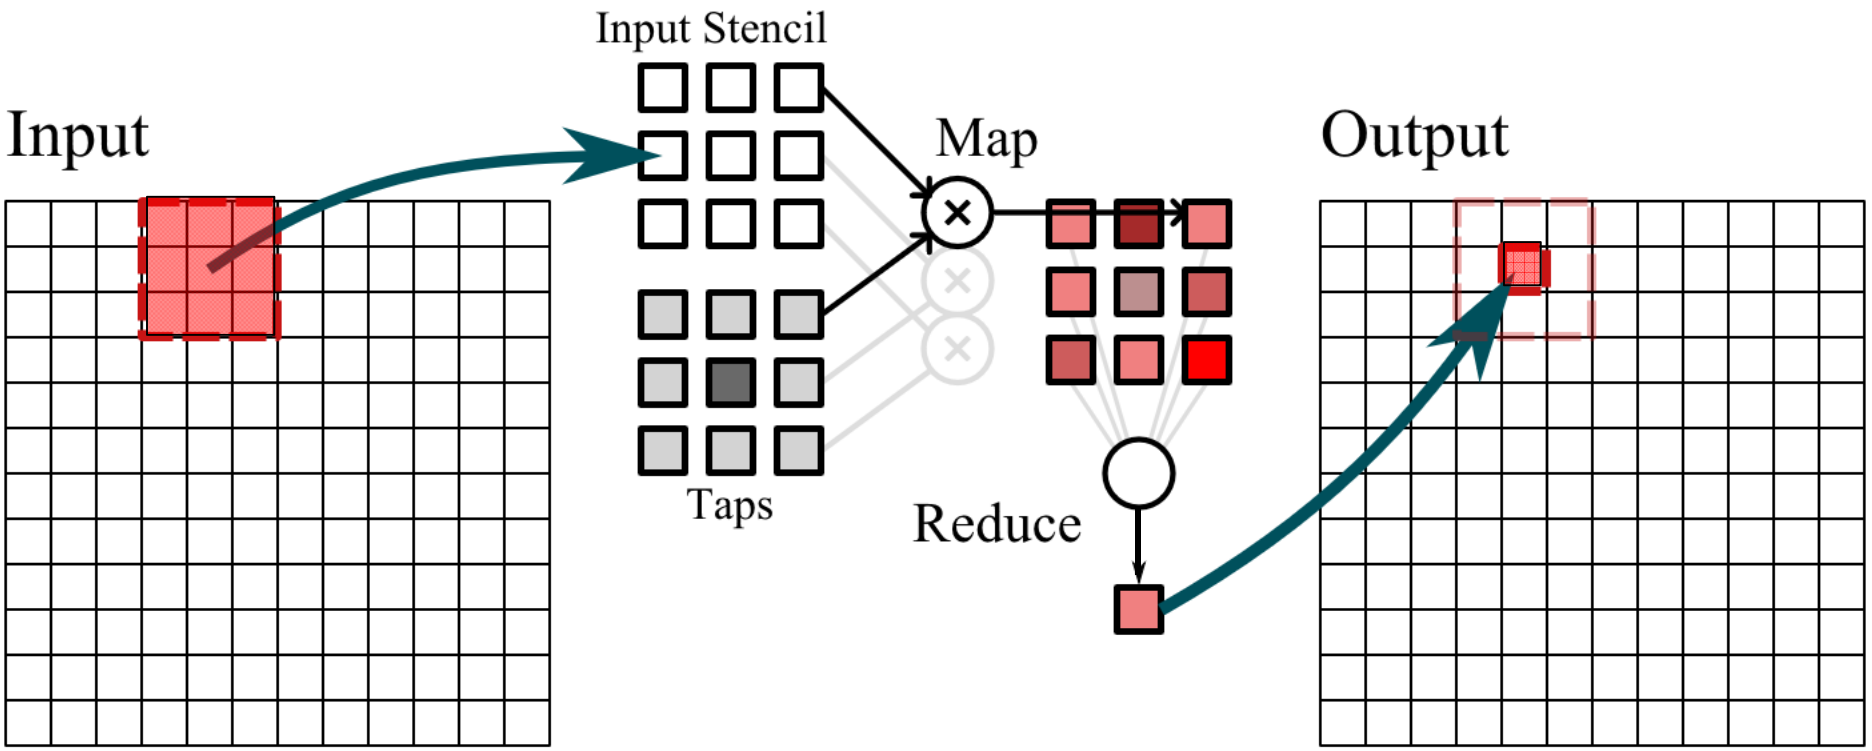

Highly Local Computation Model

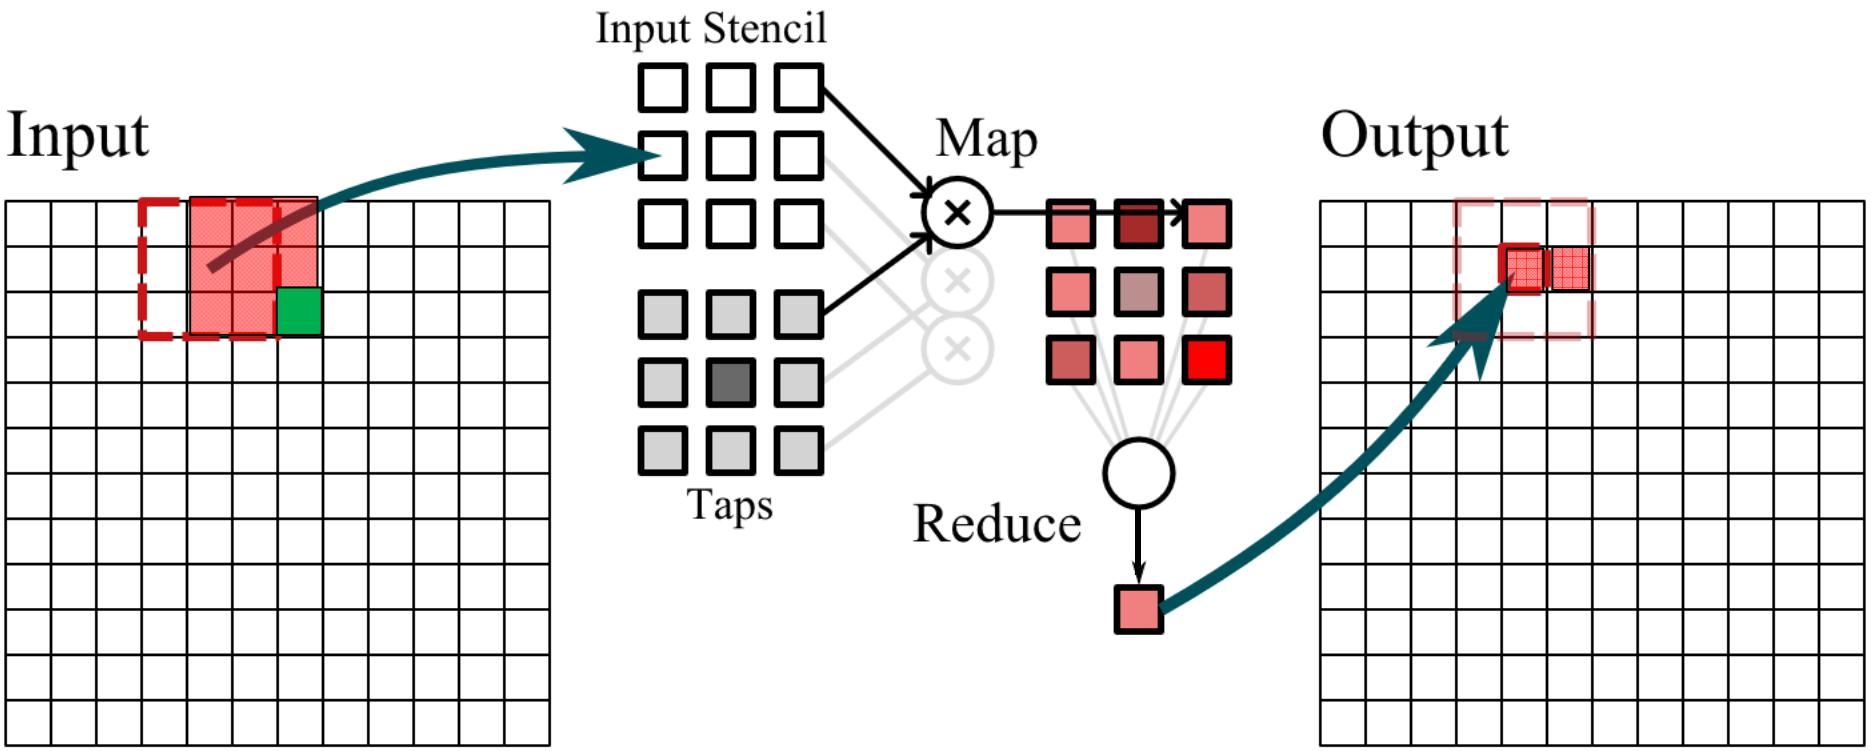

Highly Local Computation Model

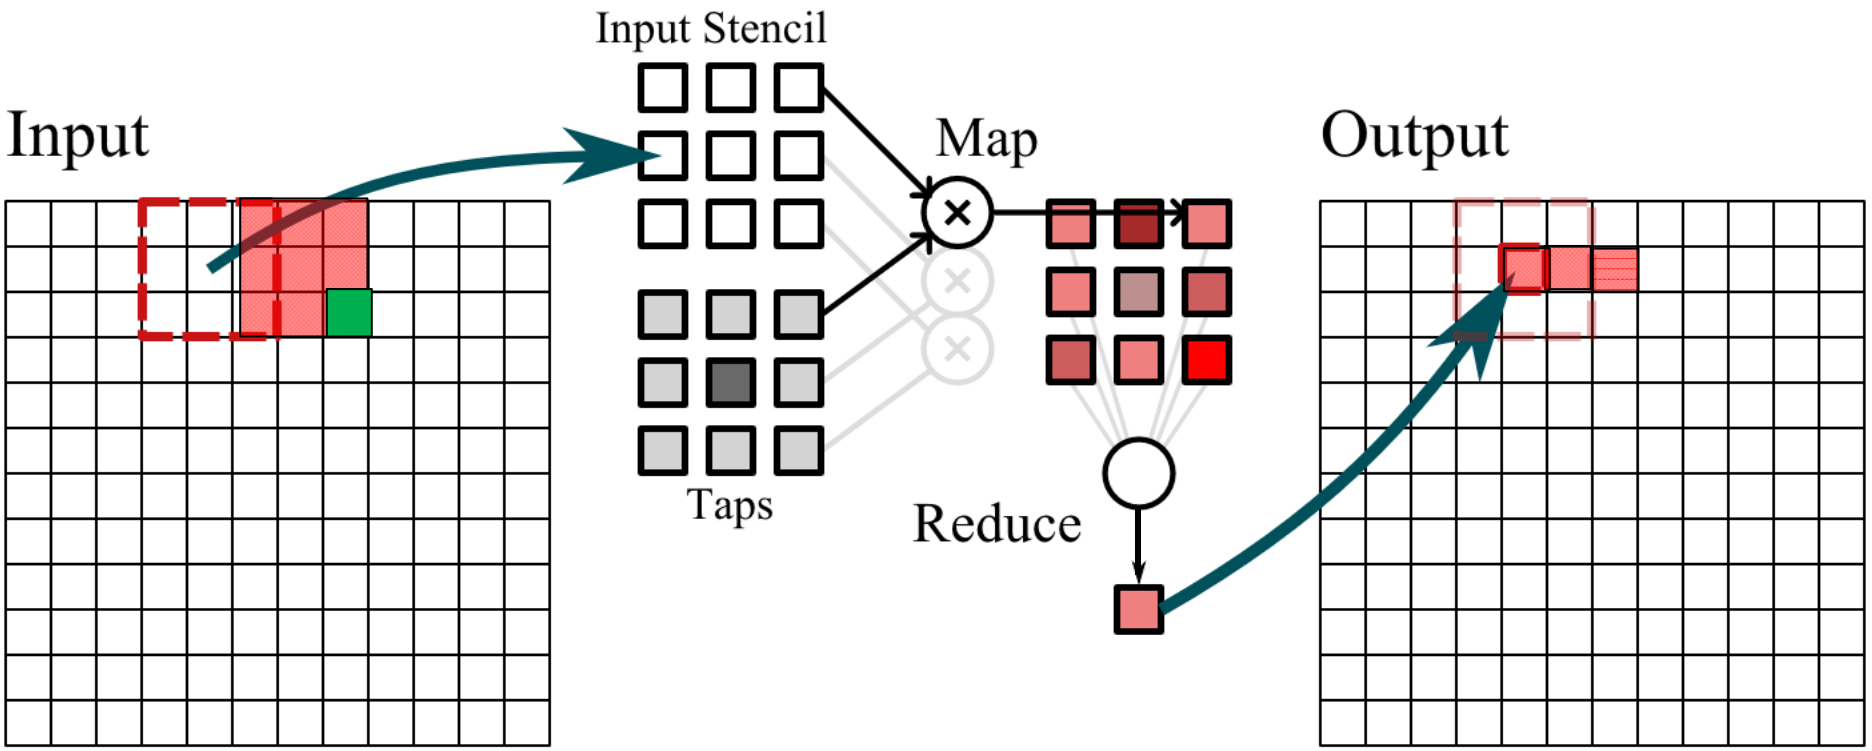

Compose These Cores into a Pipeline

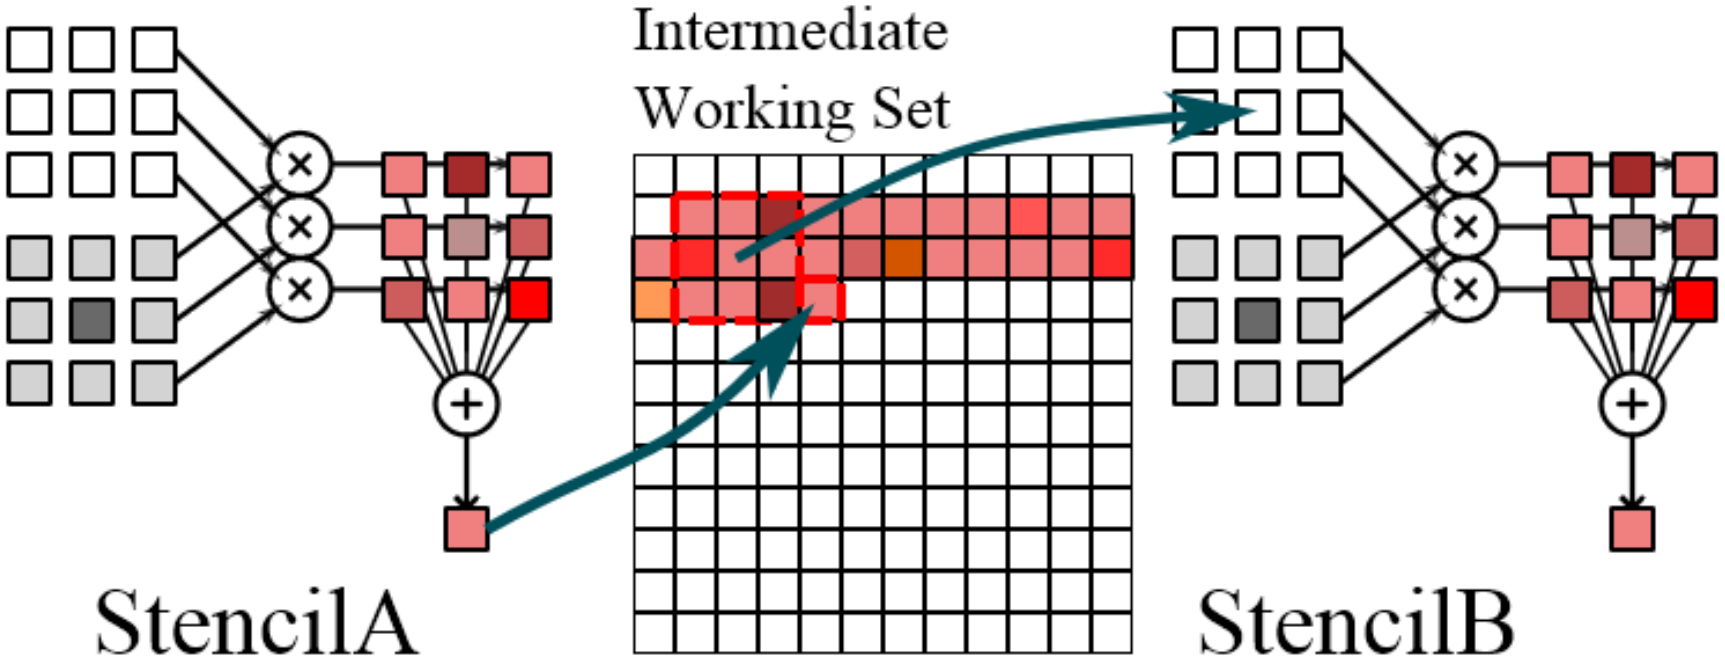

Program in space, not time

- Makes building programmable hardware more difficult

Working on System to Explore This Space

Takes high-level program

- Graph of stencil kernels

Maps to hardware level assembly

- Compute graph of operations for each kernel

Currently we map the result to:

- FPGA, custom ASIC

Enabling Innovation

You don’t just compile applications to efficiency

- Need to tweak the application to fit constraints

Need to enable application experts to play

- They know how to “cheat” and still get good results

Remember This Trade-off?

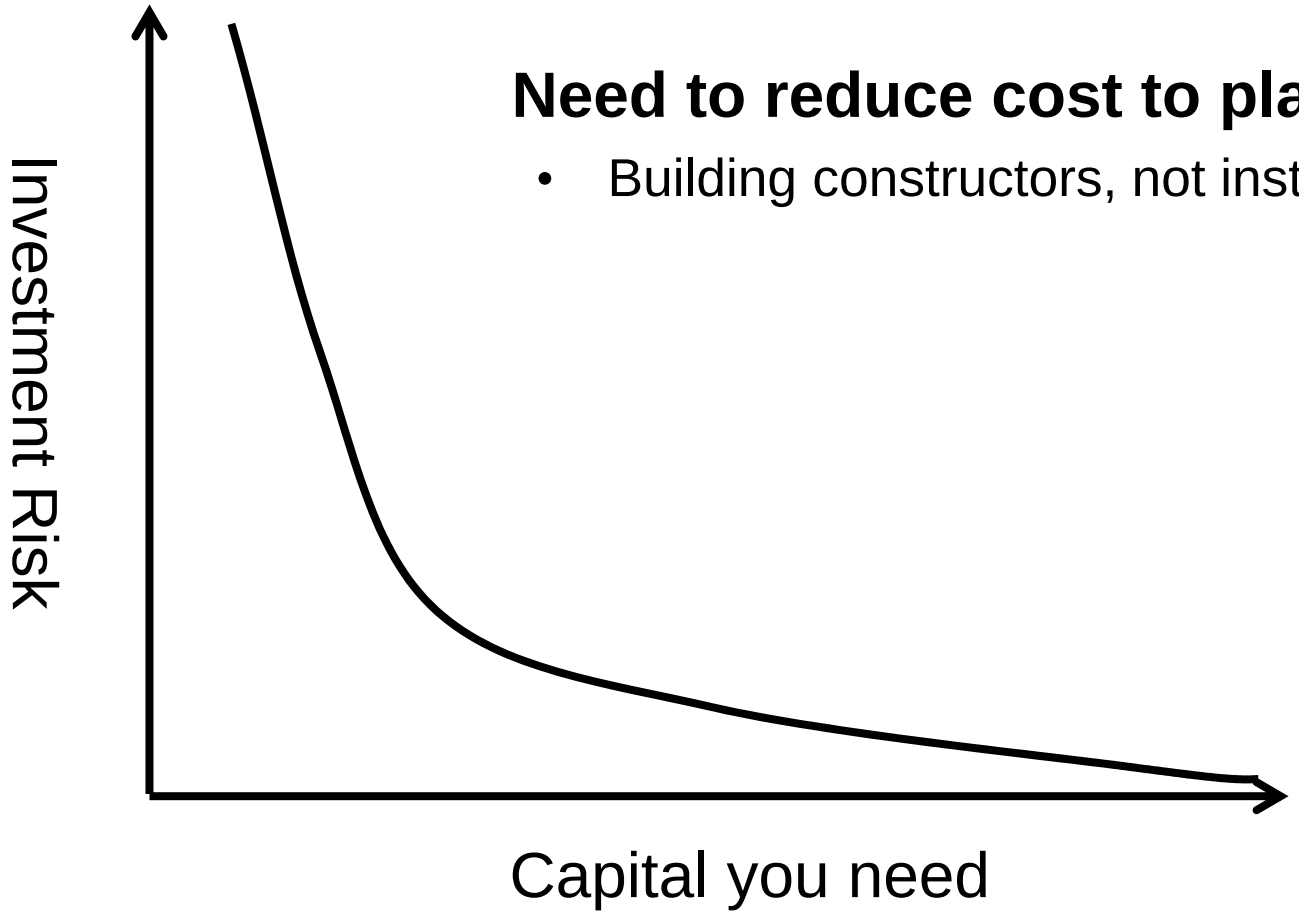

Need to reduce cost to play

- Building constructors, not instances

http://genesis2.stanford.edu/

Not All Systems Are On The Bleeding Edge

App Store For Hardware

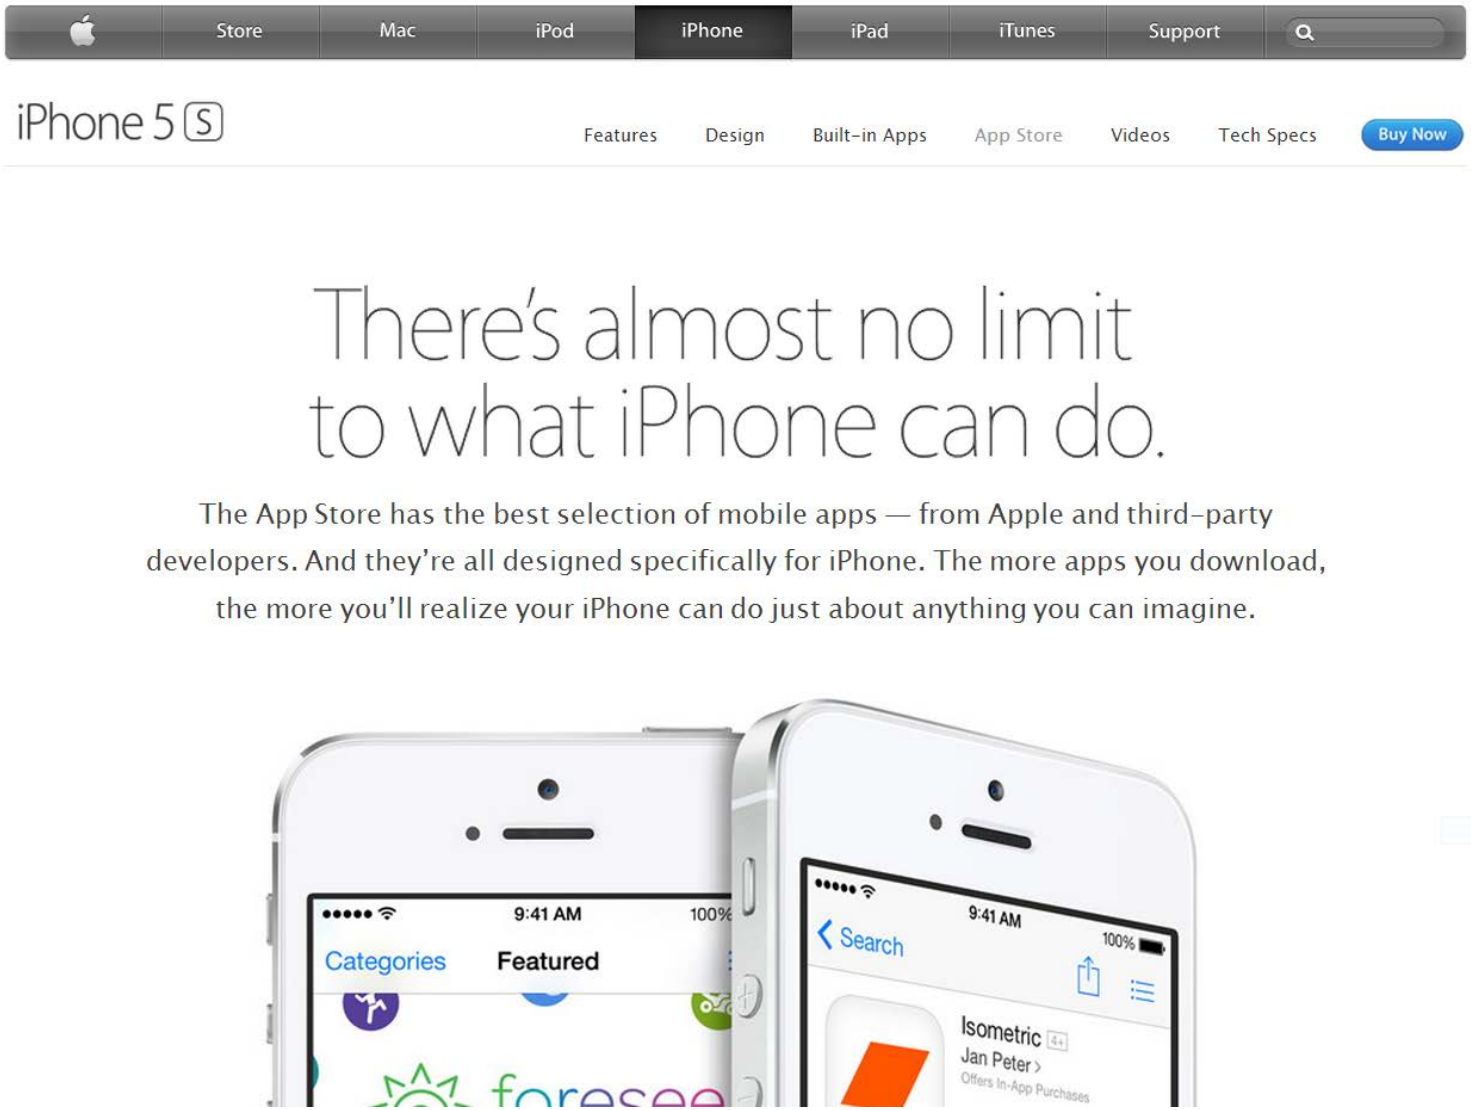

Challenge

What Arduino can do

Arduino can sense the environment by receiving input from a variety of sensors and can affect its surroundings by controlling lights, motors, and other actuators. The microcontroller on the board is programmed using the Arduino programming language (based on Wiring) and the Arduino development environment

Community

The community of Arduino enthusiasts is vast, and includes region specific groups and special interest groups. The community is an excellent further source of assistance on all topics such as accessory selection, project assistance, and ideas of all sorts.

A New Hope

If technology is scaling more slowly

- We can incorporate current design knowledge into tools
- To create extensible system constructors

If killer products are going to be application driven

- Application experts need to design them

We can leverage the 1st bullet to enable the 2nd

- To usher in a new wave of innovative computing products

Cloud 2.0 Clients and Connectivity – Technology and Challenges

Ming-Kai Tsai
Chairman and CEO

MediaTek
Hsinchu, Taiwan

ISSCC Started 60 Years Ago…

Outline

- **Reality and Macro Trend** 
- Cloud 1.0 – Smart Device Era
  - Infrastructure
  - Clients
  - Connectivity
- Cloud 2.0 – Ubiquitous Era
  - Vision of the future
  - Challenges
- Summary

Reality Check #1: Silicon Technology Scaling

"The number of transistors incorporated in a chip will approximately double every 24 months."

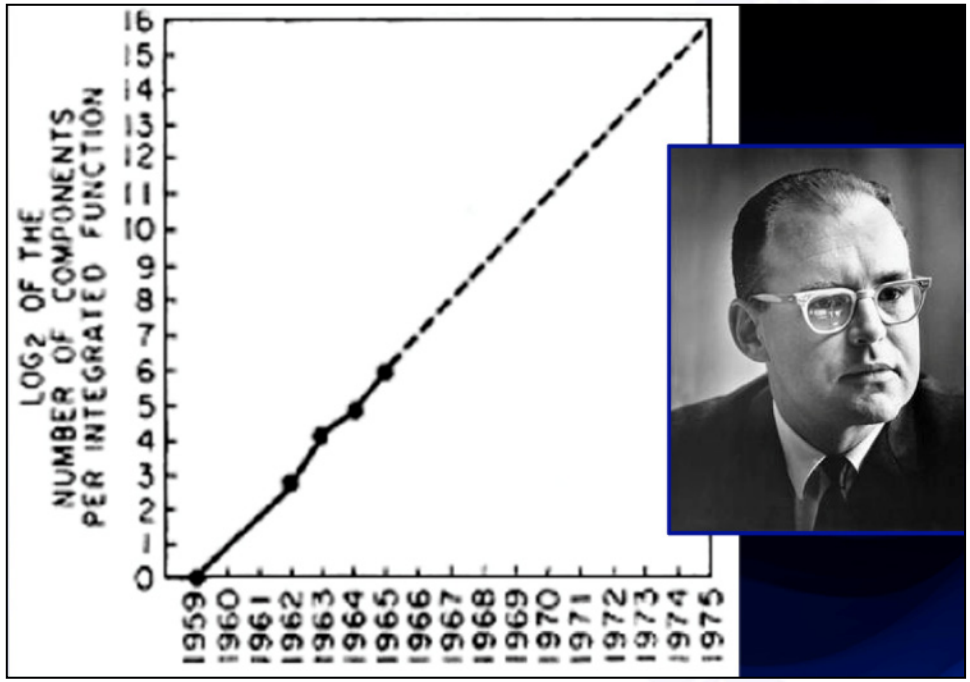

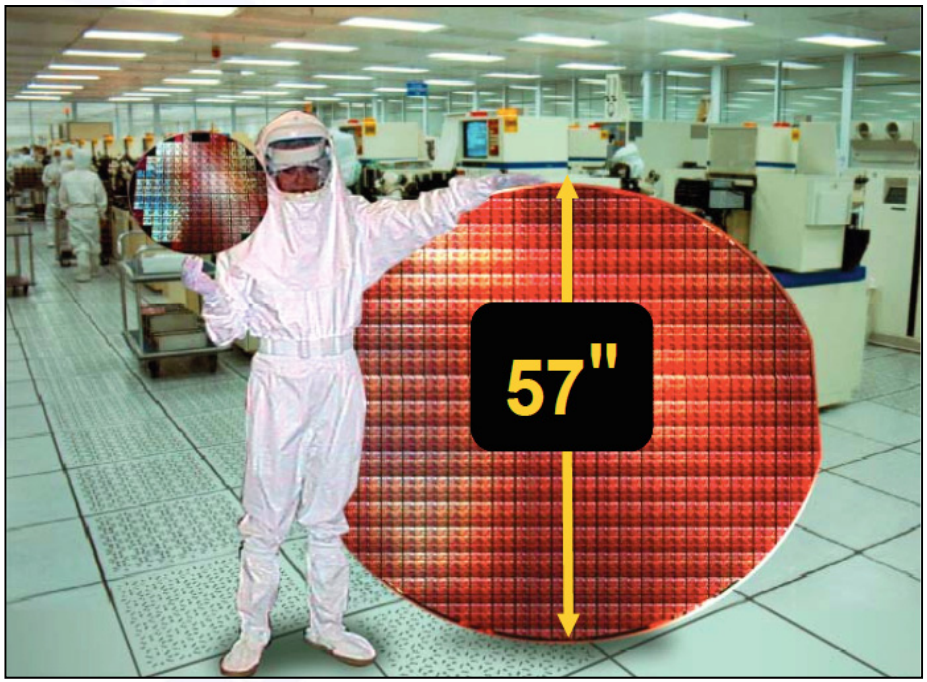

Source: 1975 circa

"2003 ISSCC: No exponential is forever… But, we can delay Forever" - Gordon Moore

Reality Check #1: Silicon Technology Scaling

"The number of transistors incorporated in a chip will approximately double every 24 months."

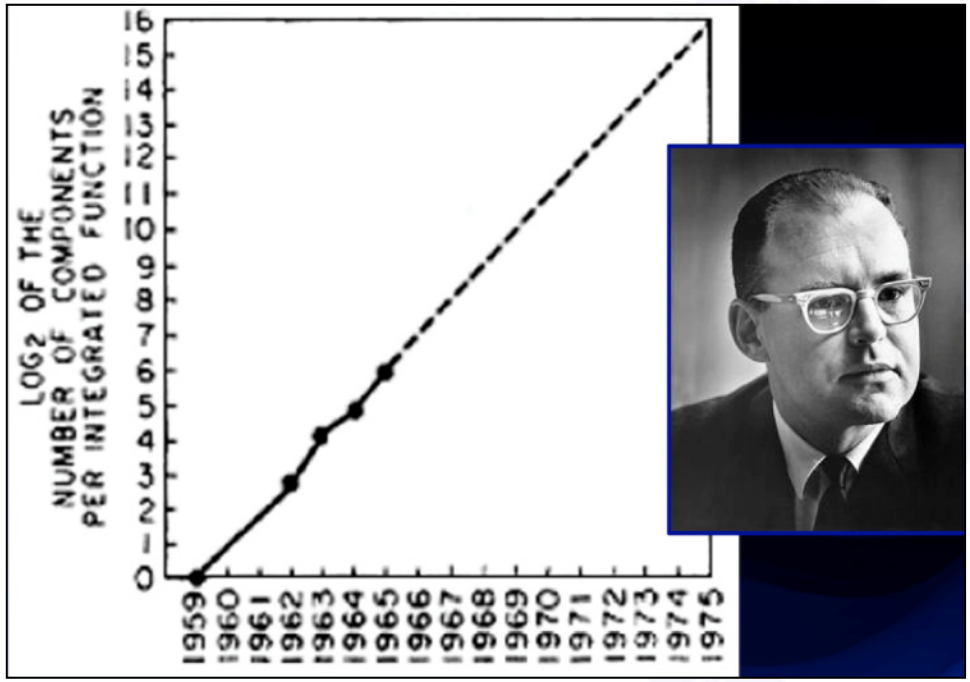

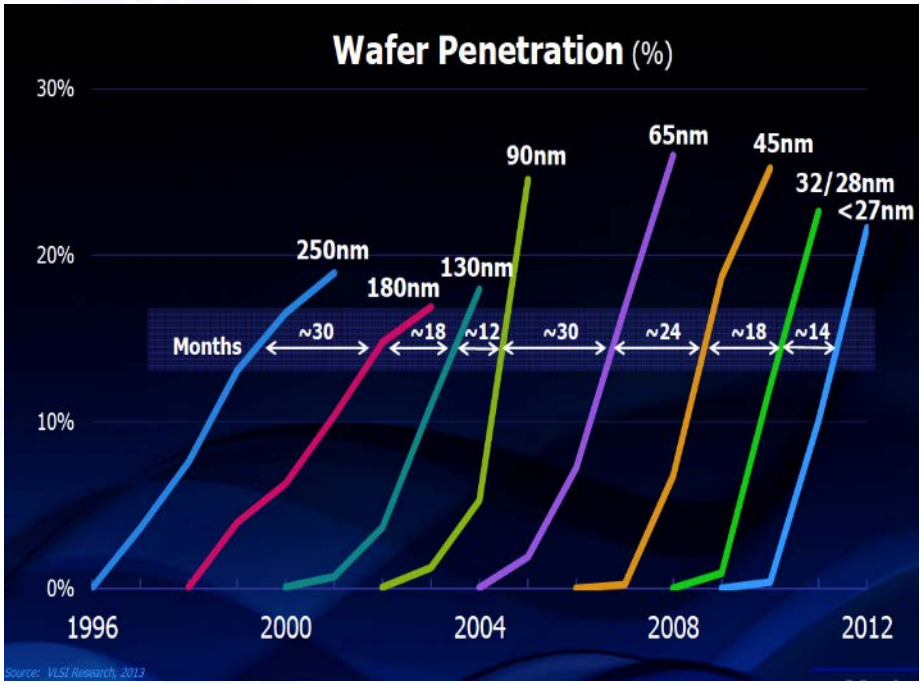

"2003 ISSCC: No exponential is forever… But, we can delay Forever" - Gordon Moore

Source: 2013 VLSI Research

While concerns on "No Exponential is Forever", technology adoption continues

Reality Check #2: Transistor Cost

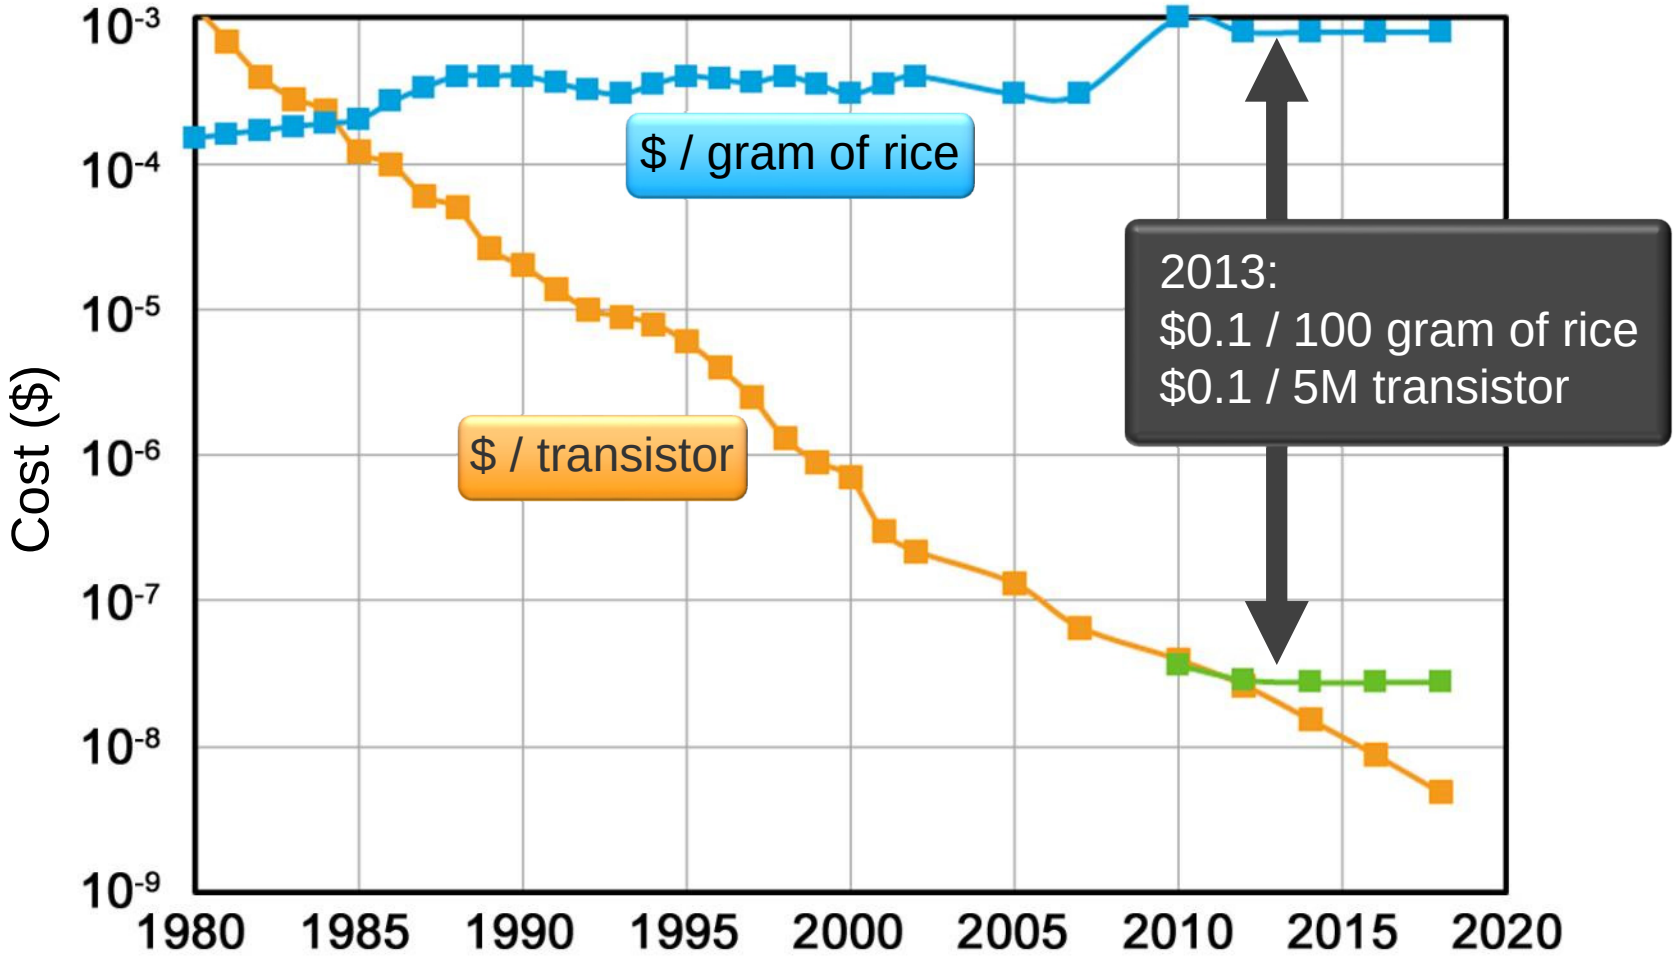

$0.1/meal of 100g rice = 100 xtors (1980) = 5M xtors (2013)

Positive Technology-Market Expansion Loop

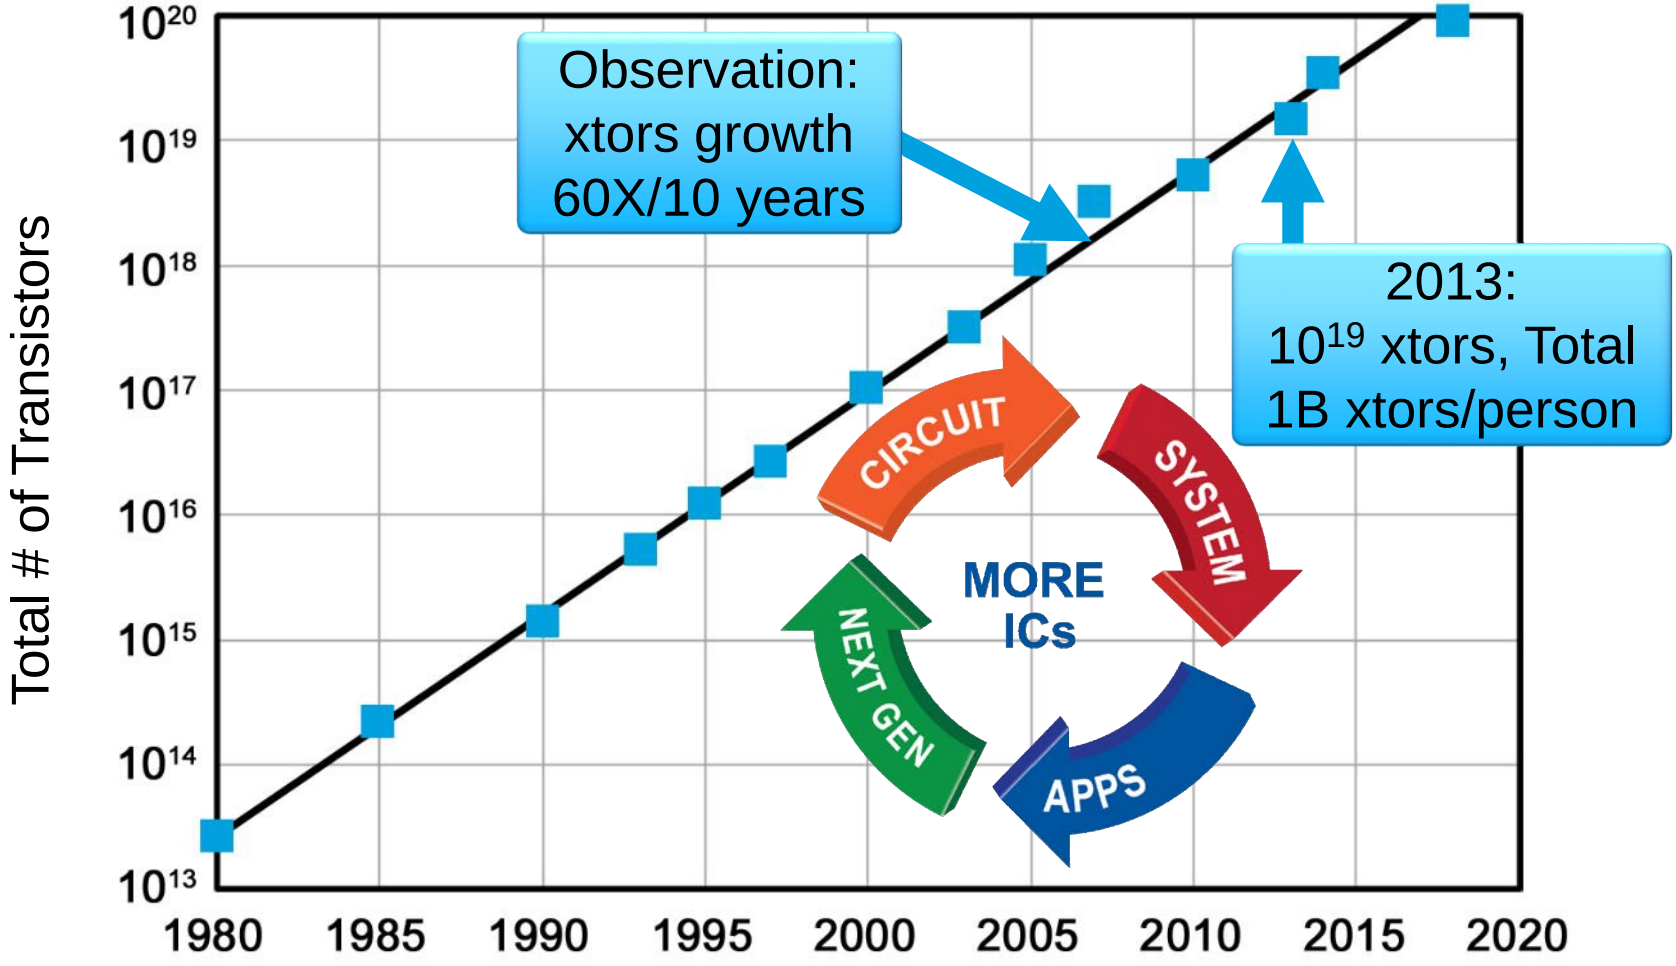

Source: 1) SEMI's WW silicon shipment statistics 2) EEMI450 3) Transistor counts

Computing Cycle Evolution

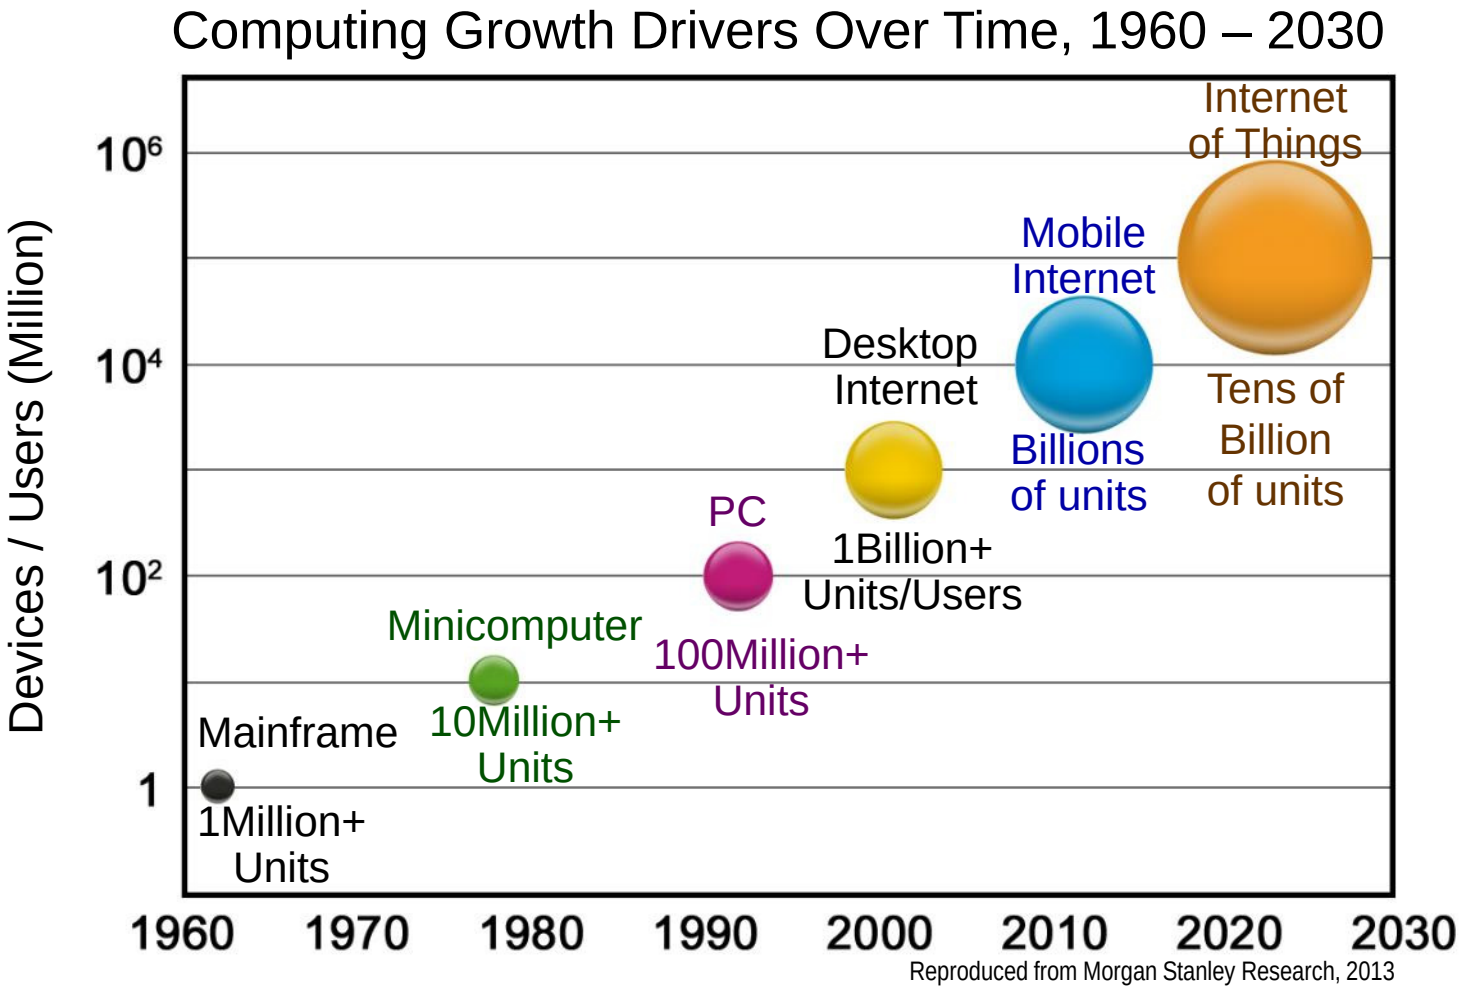

Major Technology Cycles = 10x More Users & Devices
Driven by: 1) Lower Price, 2) Improved Functionality & Services

MediaTek is Fueling the Global Smartphone Explosion

158
Countries

200+
Global & Local Brands

1300+
Smartphone projects

2013: 30% of worldwide cell phones have MediaTek inside

Mobile Impact to Emerging Countries

China: Helping farmers for real-time pricing

Pakistan: Smartphones to fight Dengue

AGRICULTURE

HEALTH CARE

EDUCATION

SERVICE

Thailand: One tablet per child

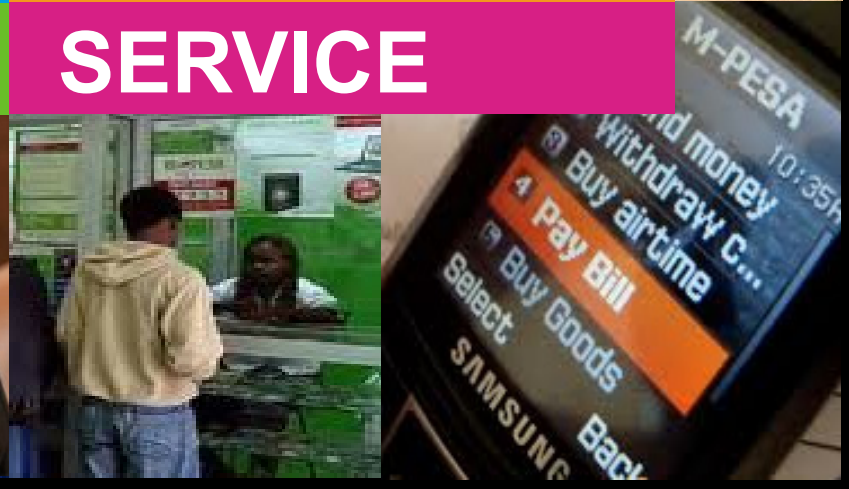

Kenya: M-PESA branch-less banking

The Emerging Middle Class

Middle Class Population (%) in different regions

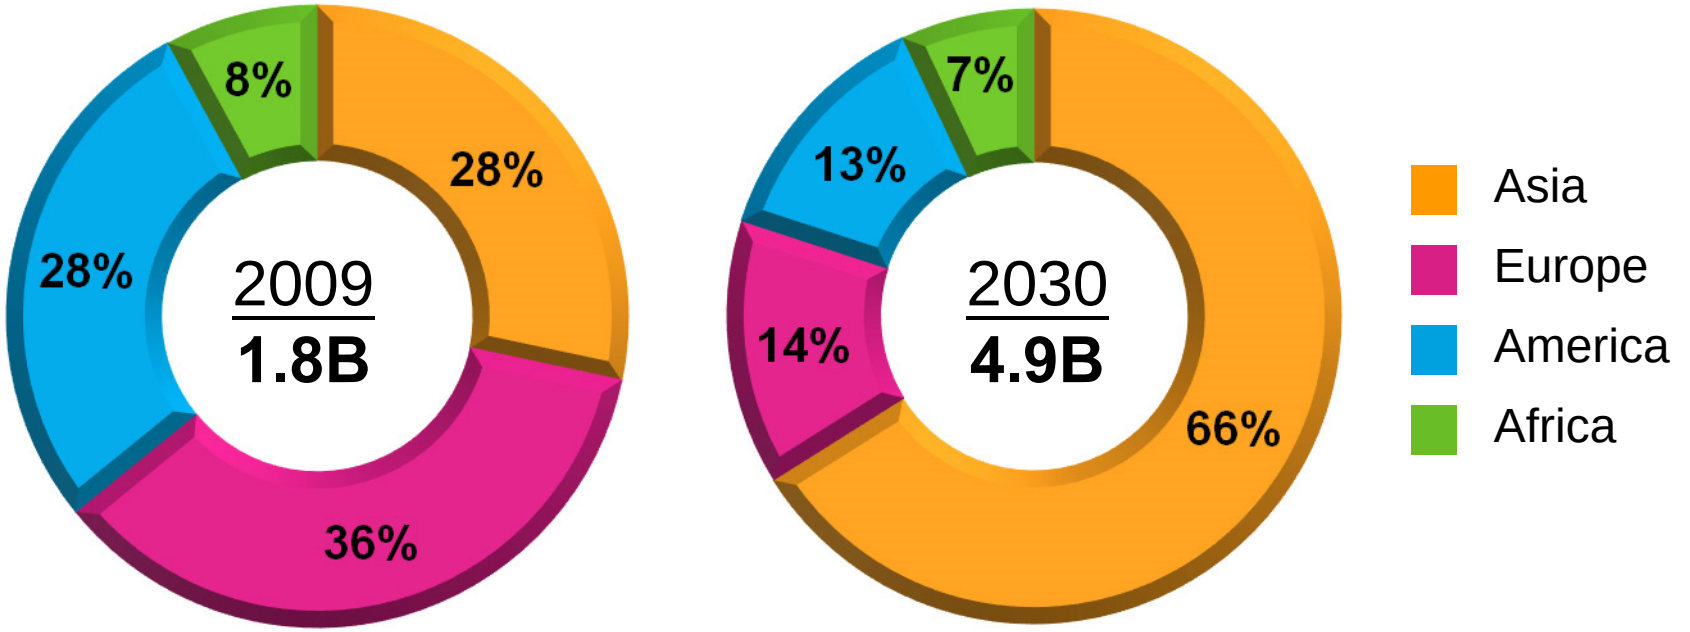

Opportunity for mobile device and commerce: $2T by 2020, on top of emerging 5B middle class by 2030

Source: 1) OECD Development Centre Working Paper No. 285, January 2010. 2) Yankee Group forecast, February 2014

Outline

- Reality and Macro Trend
- **Cloud 1.0 – Smart Device Era** 
  - Infrastructure
  - Clients
  - Connectivity
- Cloud 2.0 – Ubiquitous Era
  - Vision of the future
  - Challenges
- Summary

Cloud 1.0: Evolution to Smart Device Era

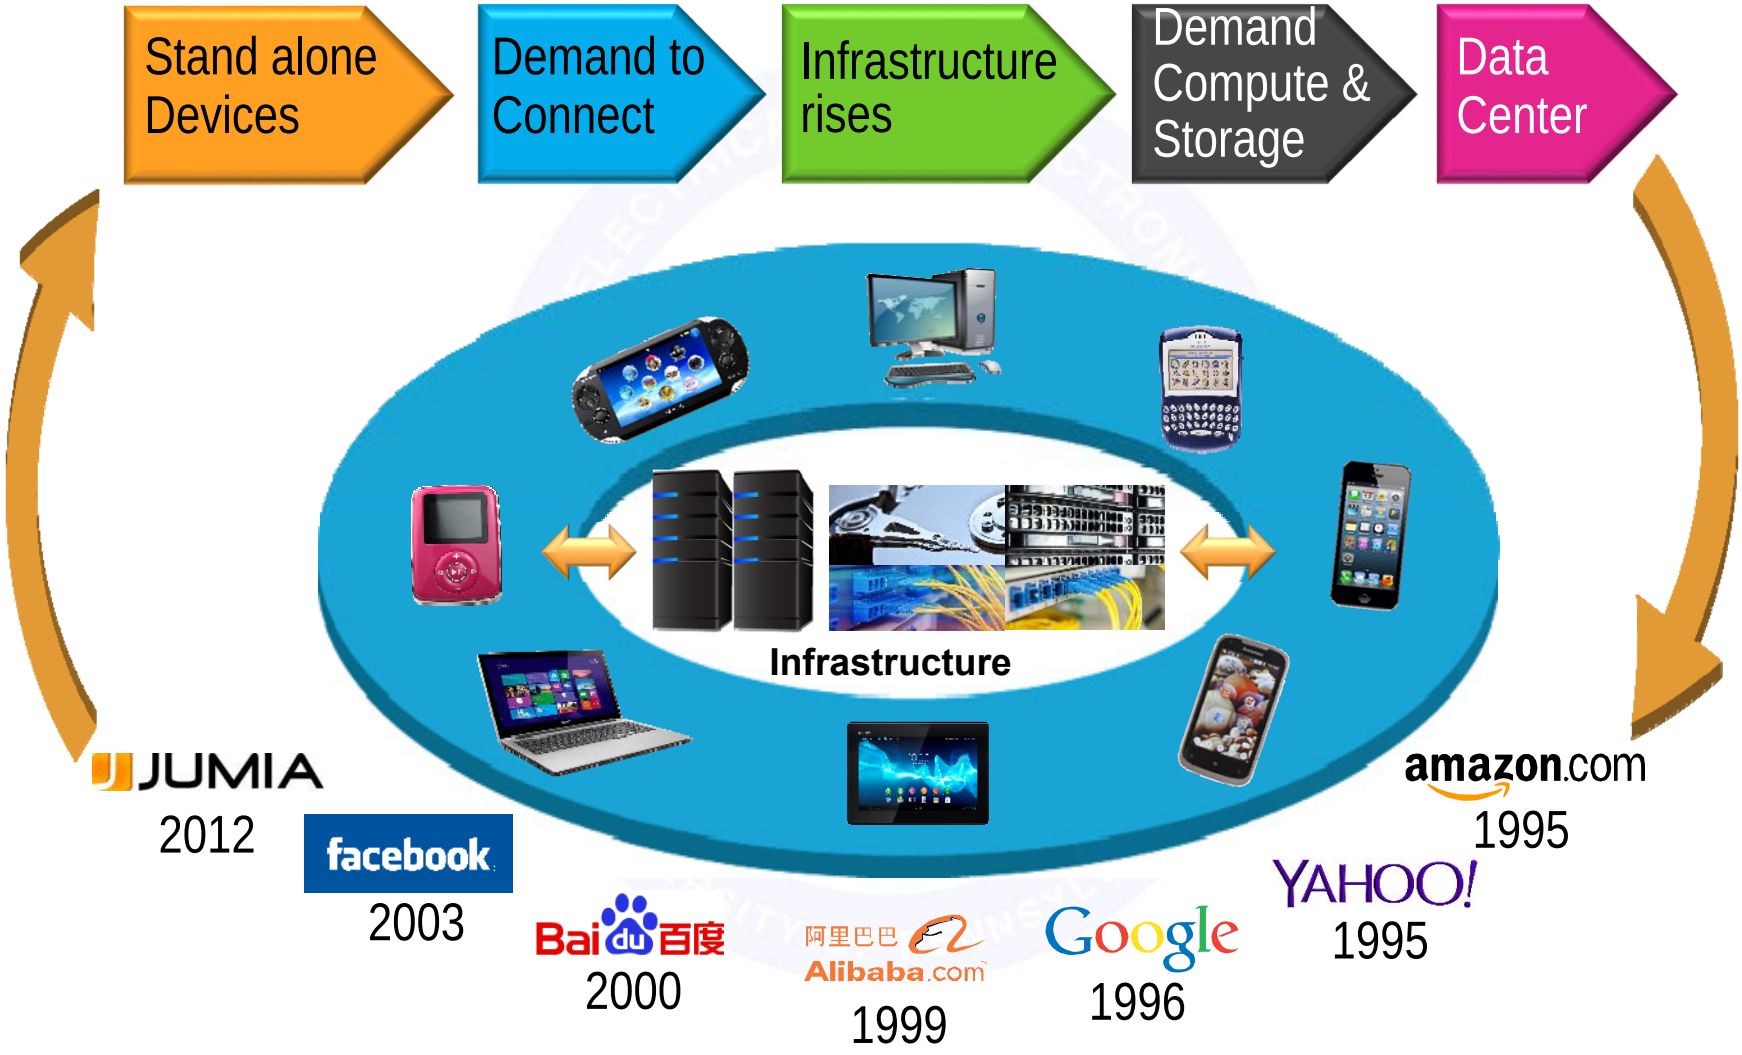

Fixed ➡ Wireless; Standalone ➡ Connected; P2P ➡ Many-2-Many

Cloud1.0: 3 Key Components

- Infrastructure
- **Client** 
  - Smartphone
  - Tablet
  - Car Electronics
  - Home Entertainment
  - Game
- Connectivity

Smartphone/tablet User-reality

Total battery usage/day

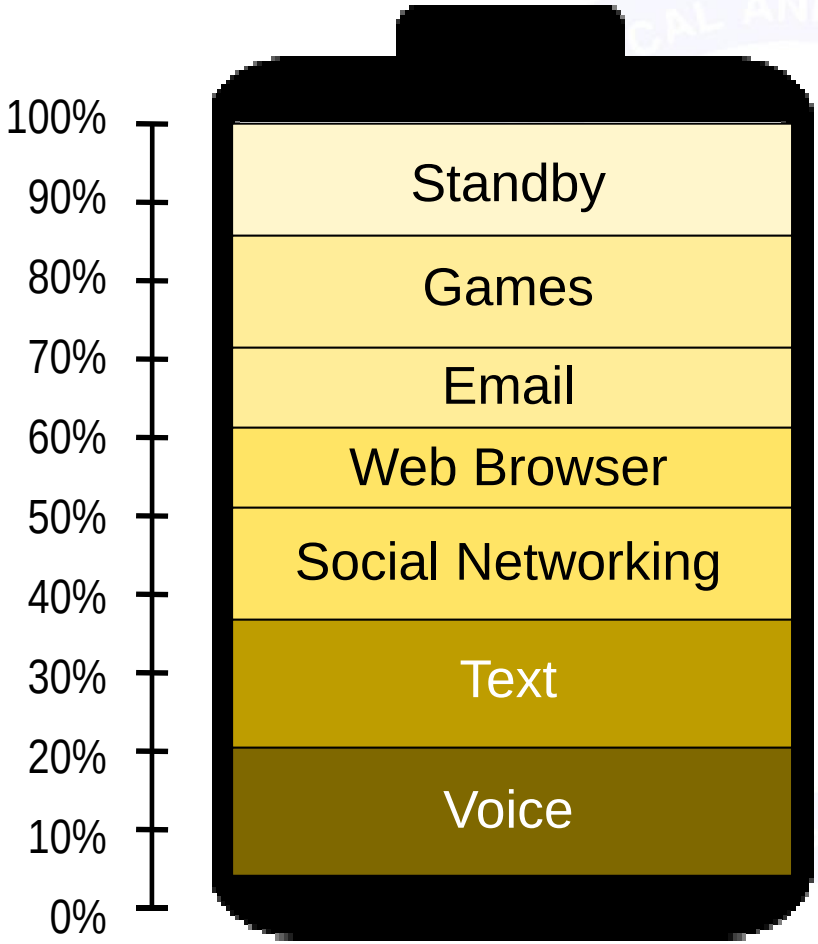

Performance:
Gaming, Web browsing

Power:
All use cases

For best user experience, need to optimize both power and performance

CPU Performance Trends

~3GHz wall where thermal cost became high

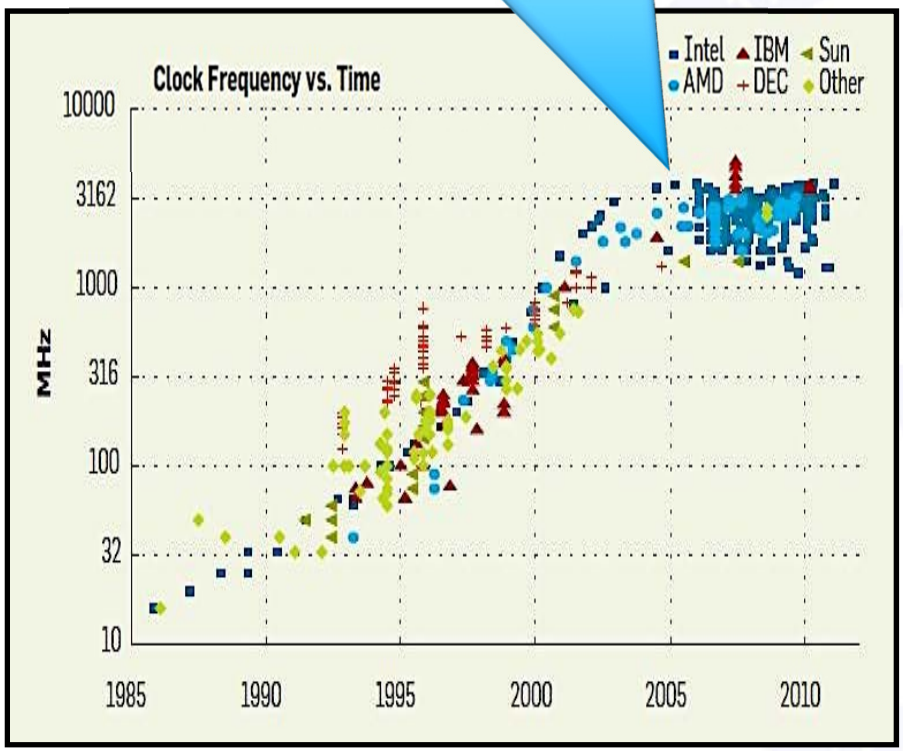

Source: A. Danowitz *et al. Communications of the ACM, April, 2012*

Battery technology not improving: energy gap

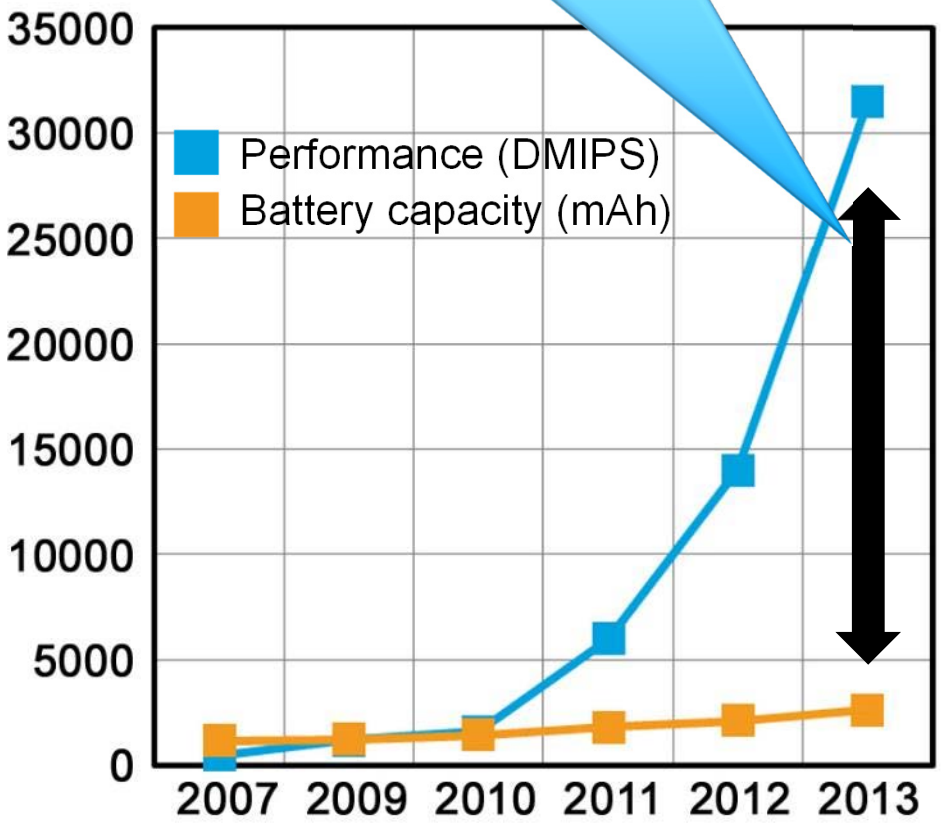

Need CKT/Architecture/System solutions to close the energy gap

Multi-core Architecture and Software Optimization

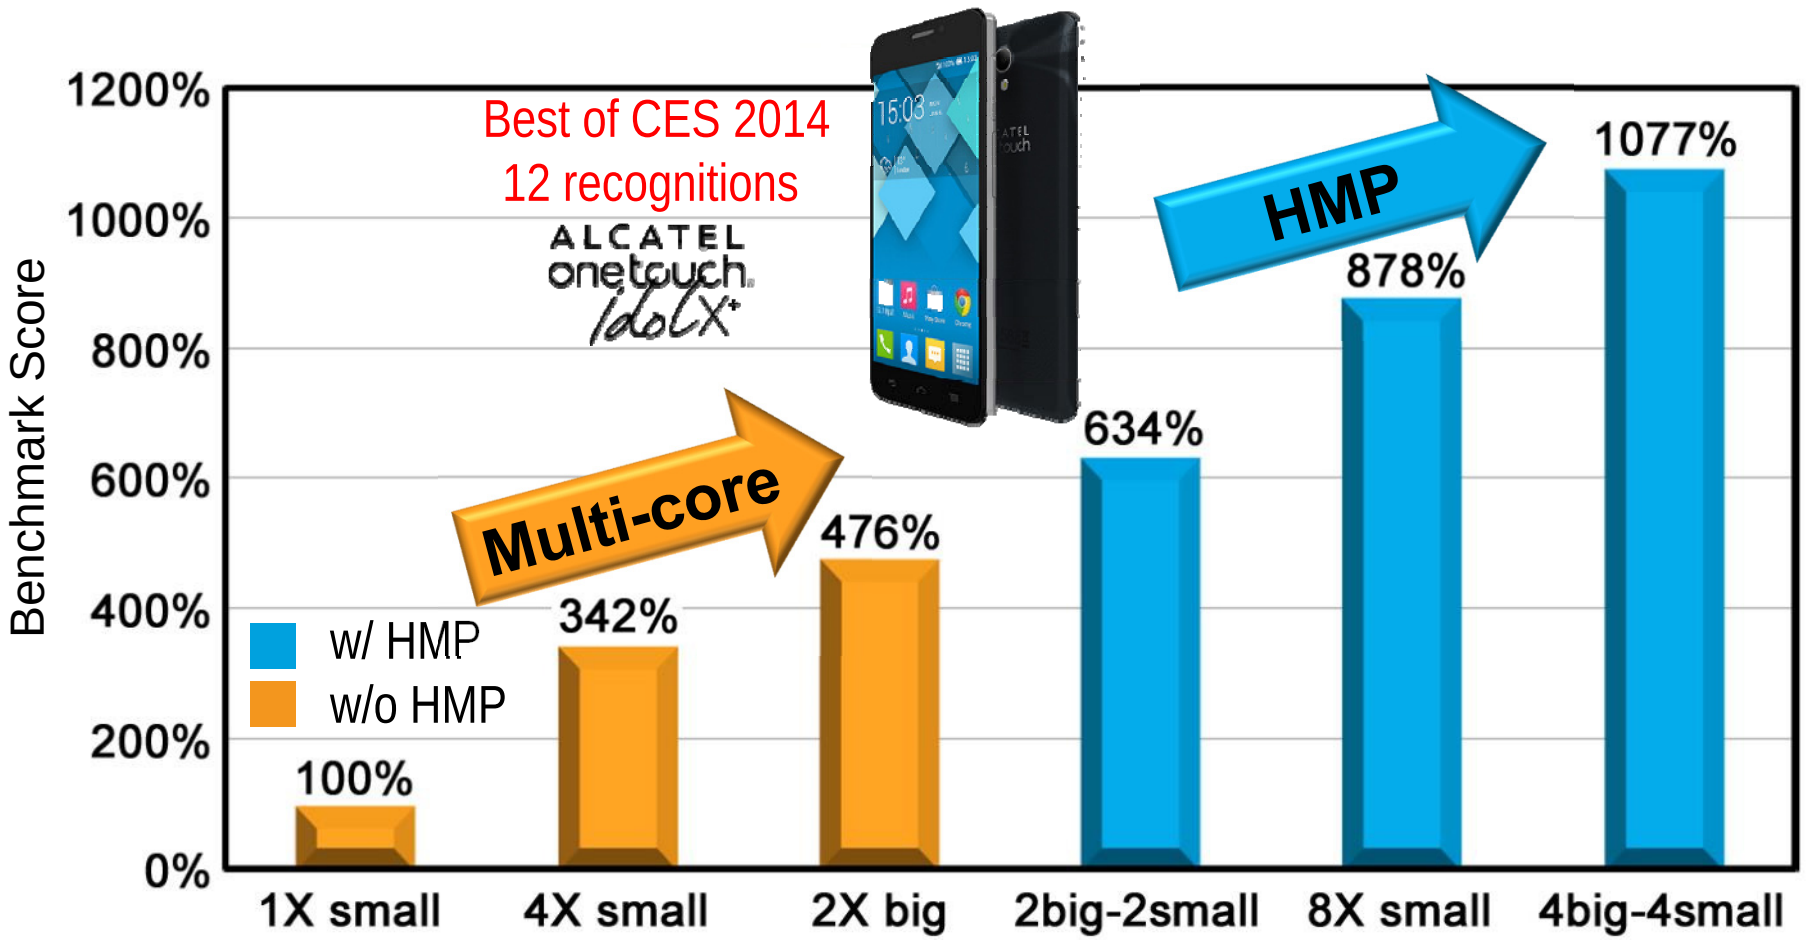

World's 1st true HMP, octa core "Best of CES 2014: 12 Awards"

HMP: Heterogeneous Multi-Processing

Performance, Thermal and Power (PTP) Technology for Lowest power

Paper 10.3
ISSCC 2014

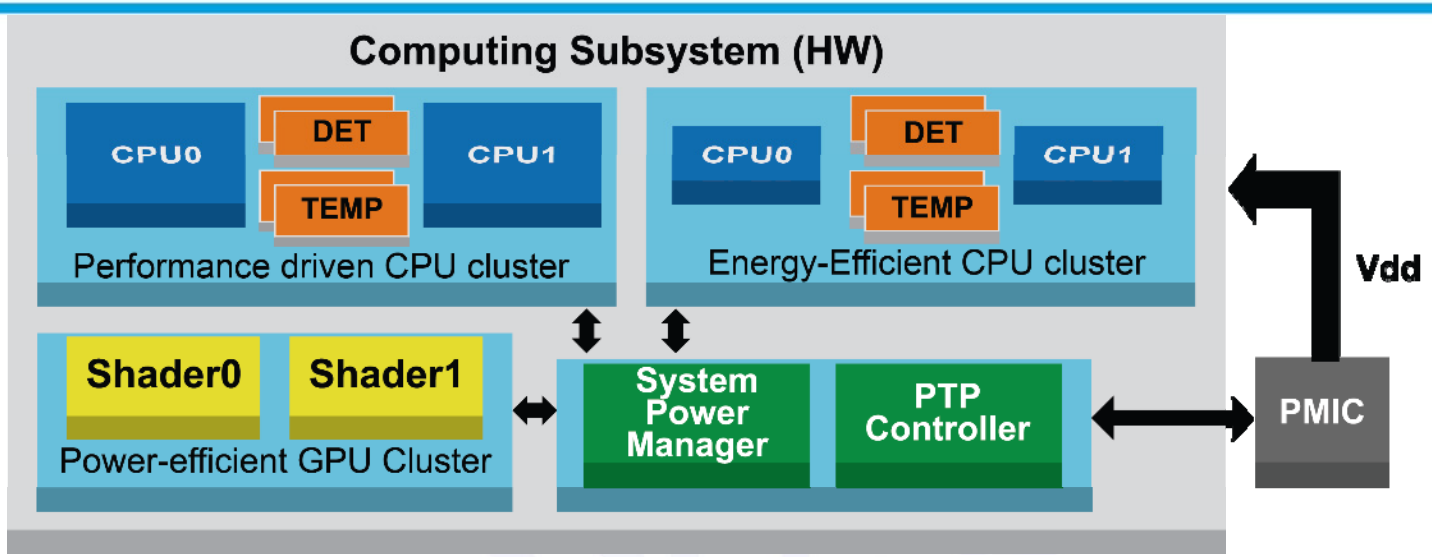

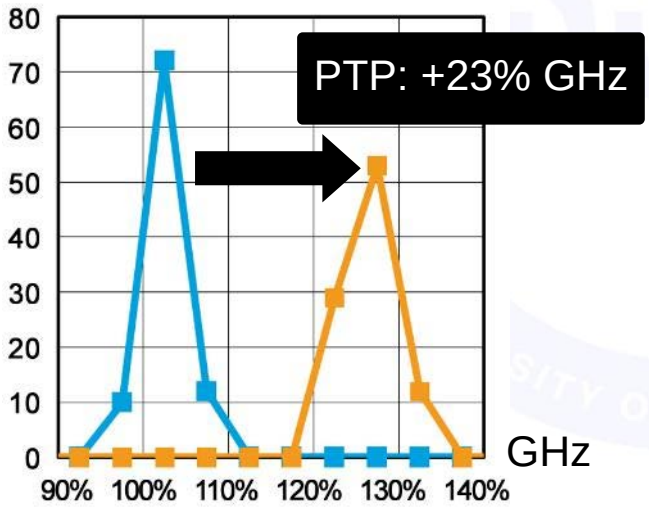

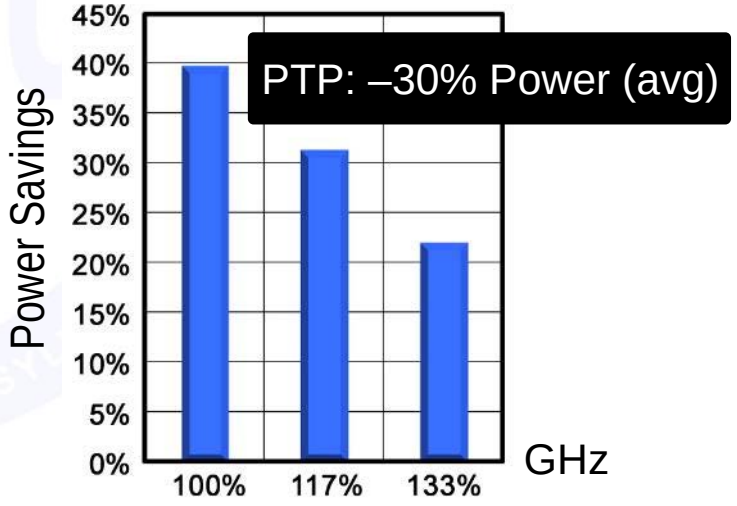

Industry first true quad-core with HMP and PTP technology for 30% lower power and 23% more performance

Cloud1.0: 3 Key Components

- Infrastructure
- Client
- **Connectivity** 
  - Wi-Fi: 802.11 a/b/n/ac/ad/…
  - Cellular: 2G, 3G, LTE, …
  - BT, BLE, NFC, BAN, …
  - GPS, FM, …

Cloud 1.0: Client Communication Evolution

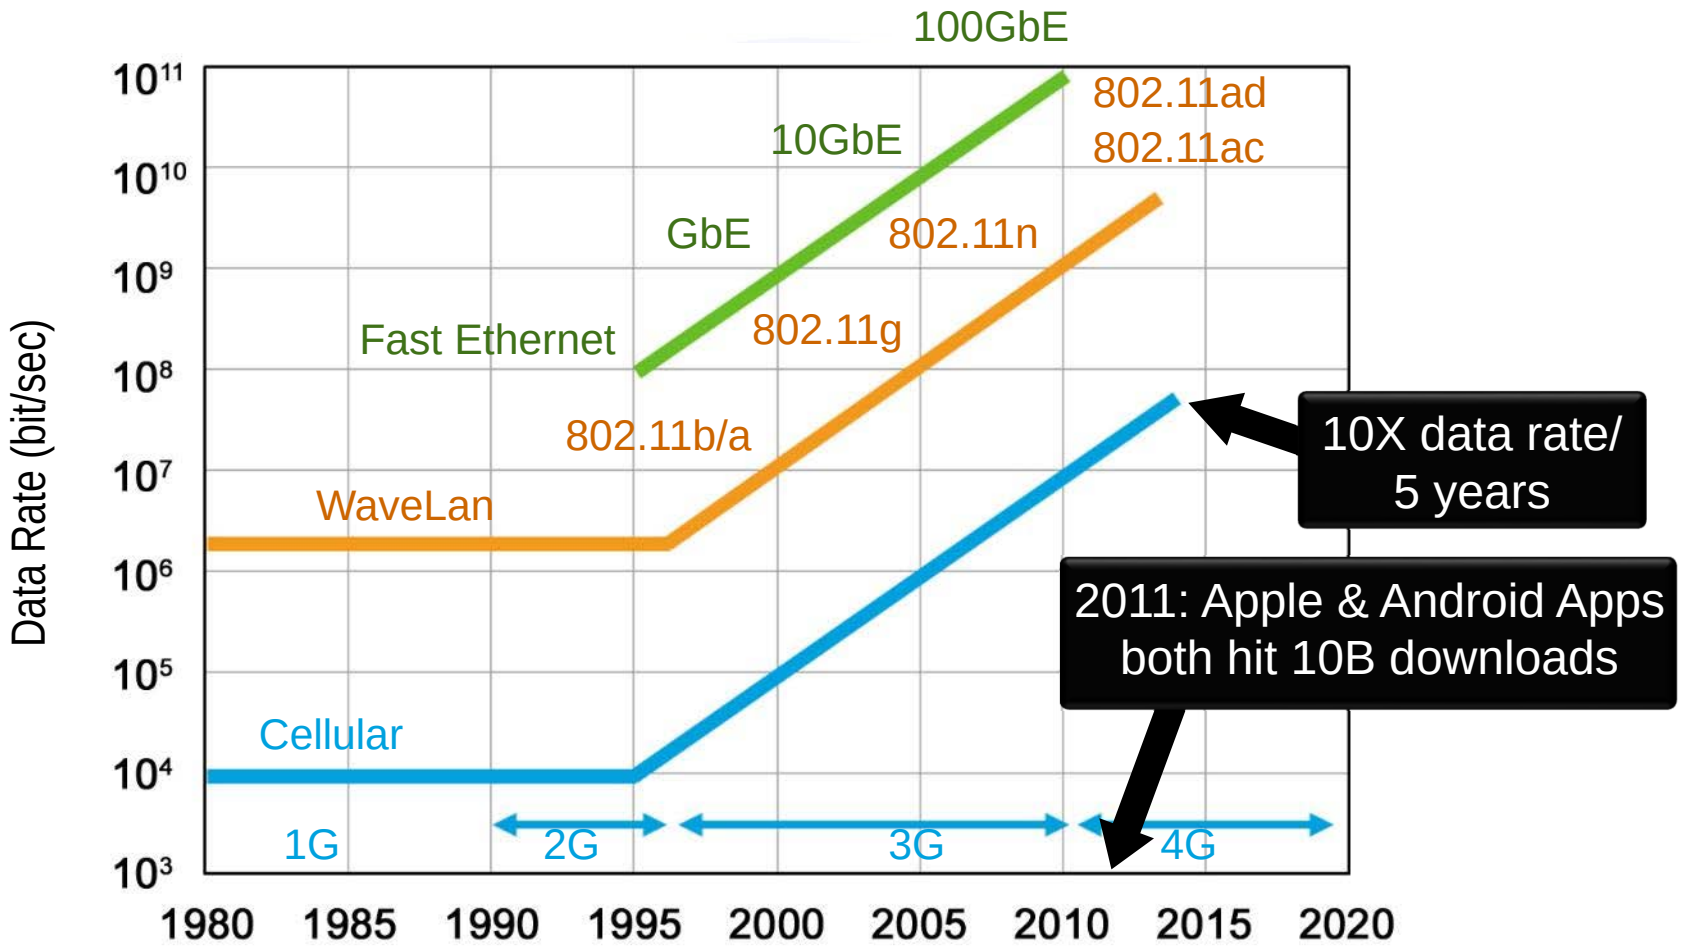

Connectivity enabled exponential data rate, more smart devices and Apps.

Mobile Technology vs. User Experience

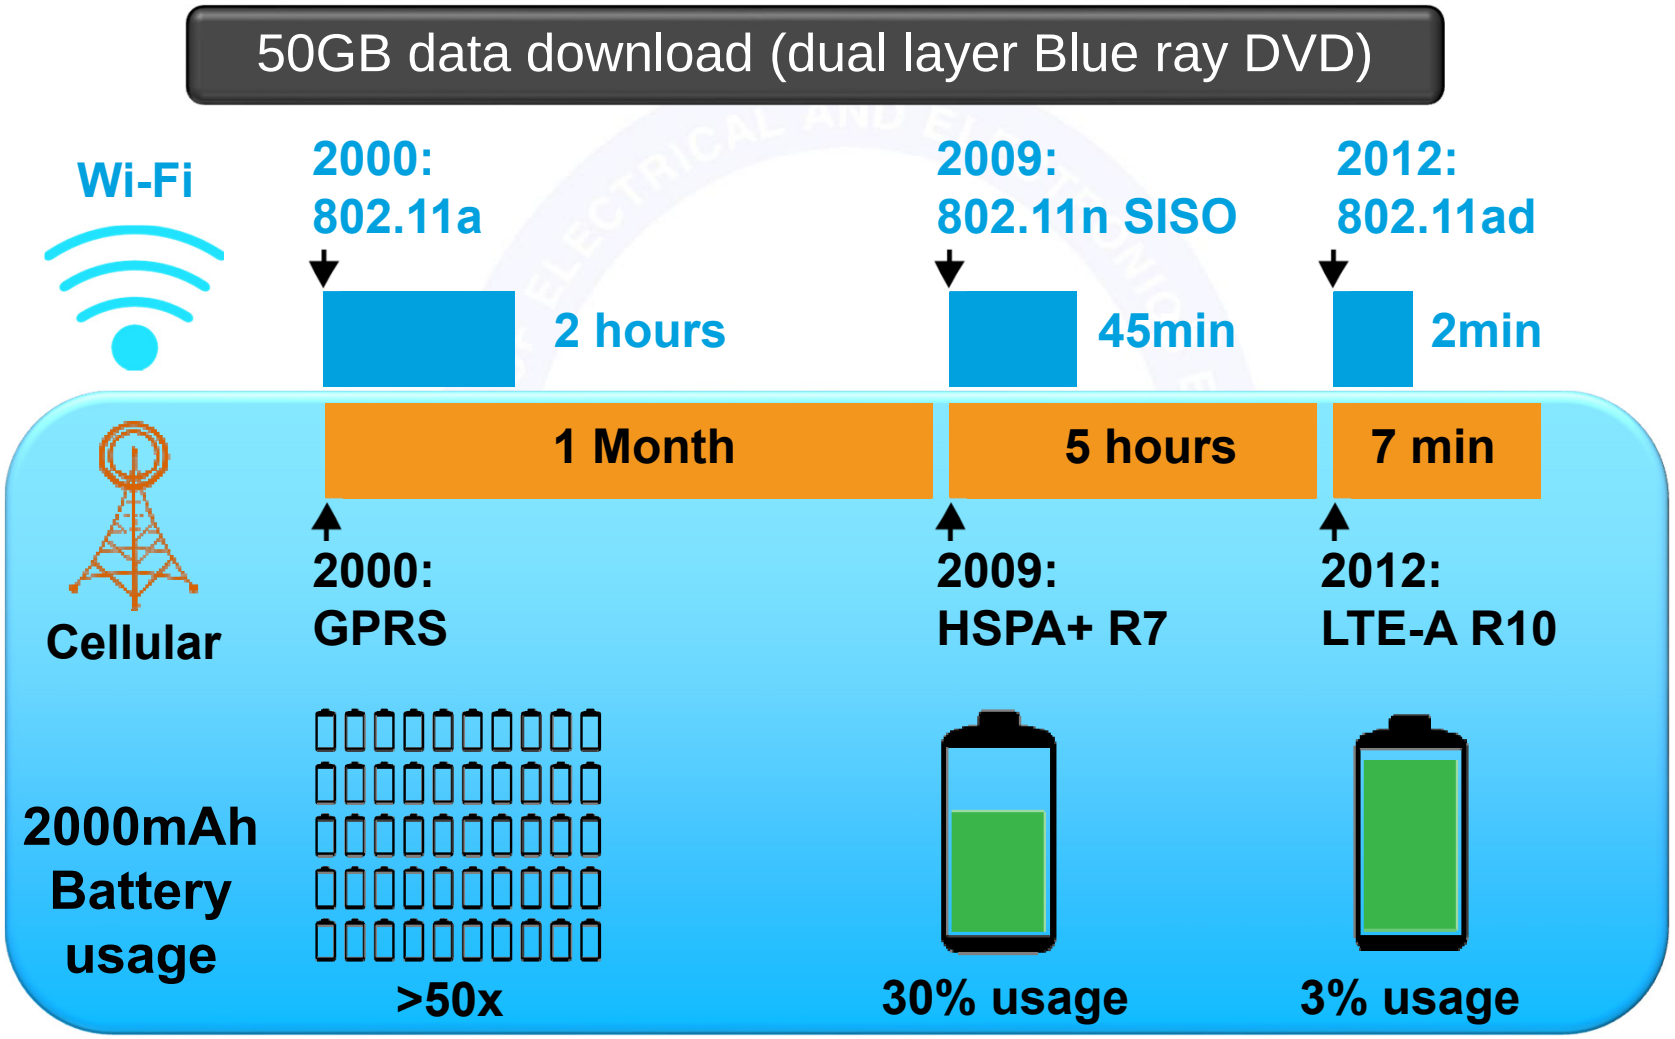

Many apps previously impossible are now part of daily life

Outline

- Reality and Macro Trend
- Cloud 1.0 – Smart Device Era
  - Infrastructure
  - Clients
  - Connectivity
- **Cloud 2.0 – Ubiquitous Era** 
  - Vision of the future
  - Challenges
- Summary

More Radio Links

TODAY:
10 Radios/Person
70 Billion Radios

2030:
100 Radios/Person
800 Billion Radios

$100B+ Silicon opportunity

Cloud 2.0: Work, Home, and Personal

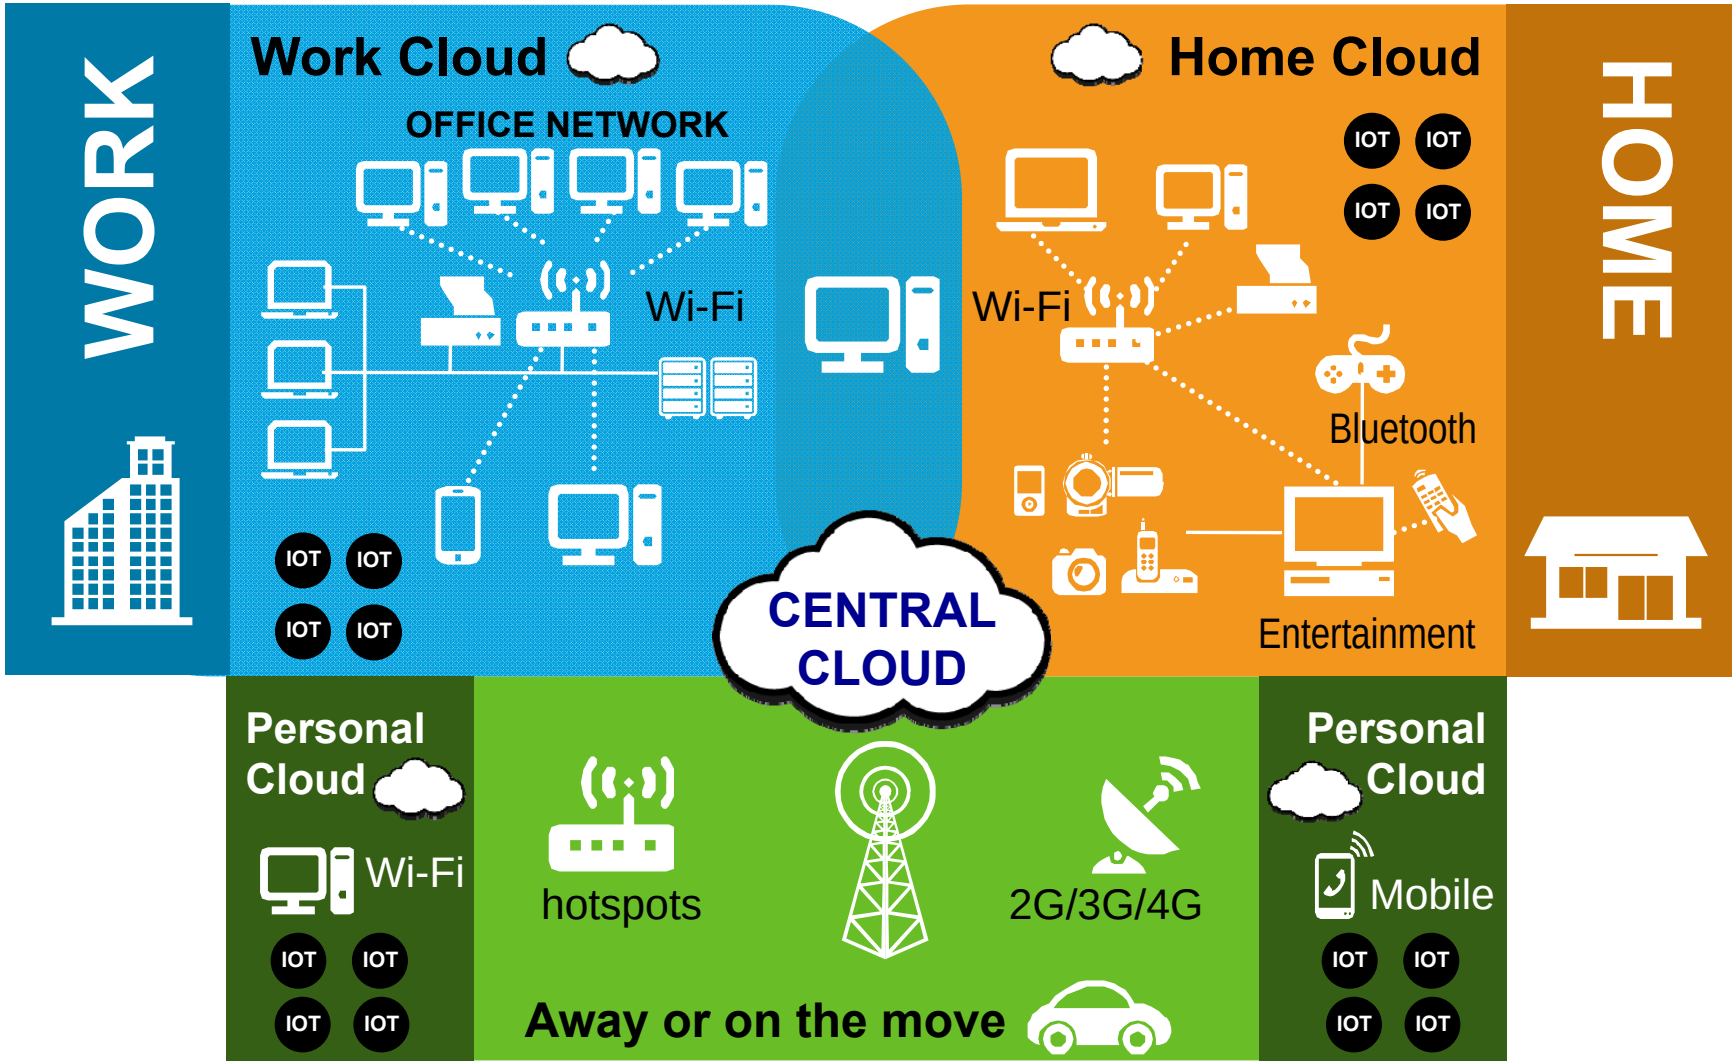

Central cloud ➡ work/home/personal cloud for ubiquitous coverage

Context-Aware Smart Devices

LBS Social

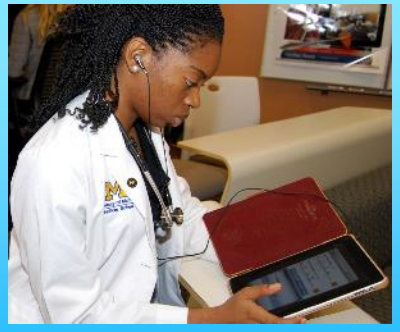

Smart Health

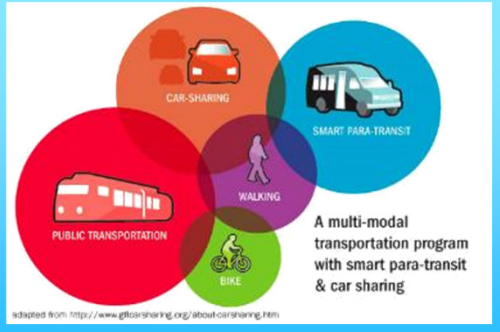

Transportation

Context-aware smart device to provide true situation-adaptive user experience

Working Toward Cloud 2.0

Wireless Charger

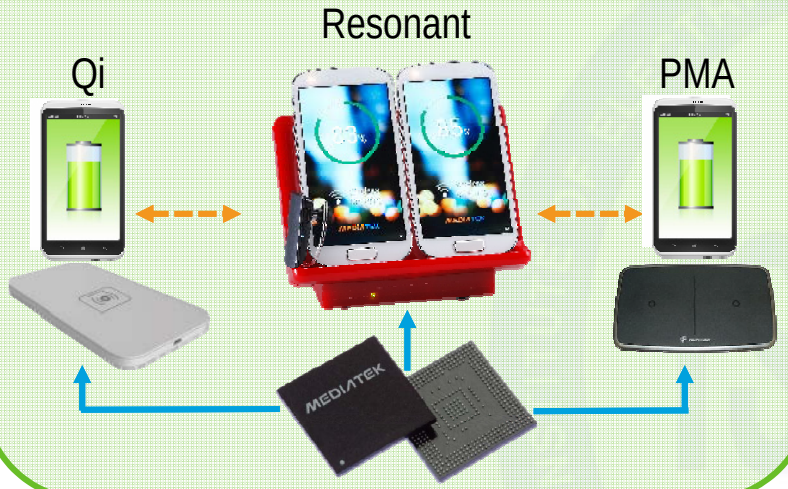

Aster: SoC for Wearable

Best Mobile Processor MT6572
2013 Analysts' Choice Awards, The Linley Group

HSA: Harmonizing the Industry Around Heterogeneous Computing

HSA FOUNDATION


Natural UI Gesture


Augmented Reality


Content Everywhere


AV Content Management

Dream #1: Personal Digital Assistant

Personal Daily Briefing Assistant PLUS **True 5 Senses Multimedia Experience**

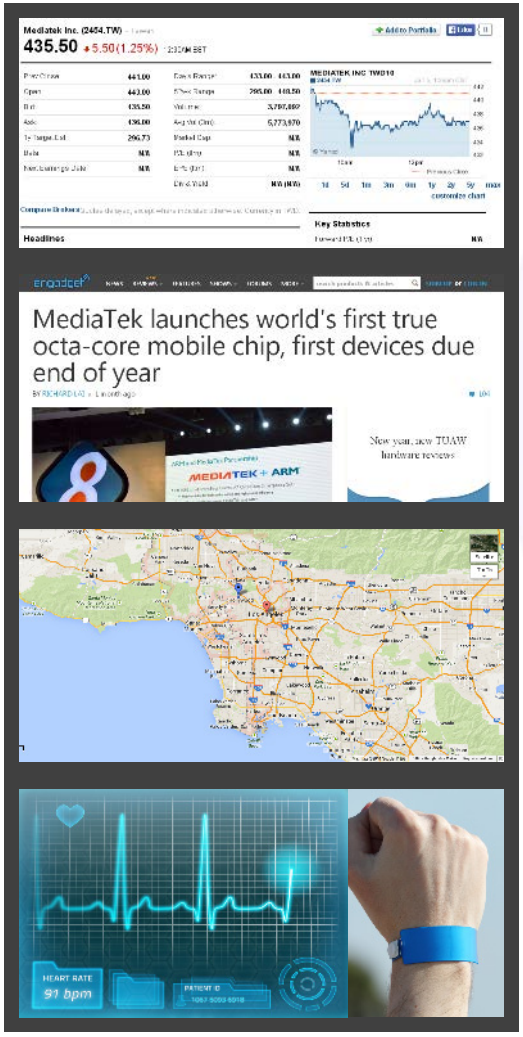

Gather personal financial data

Collect relevant news articles and technical reports

Read to me as I am driving to work

My health vs. fitness goals

Taste & Smell

Touch & Feel

Sight & Hearing

Dream #2: Smart Devices for Emerging Countries

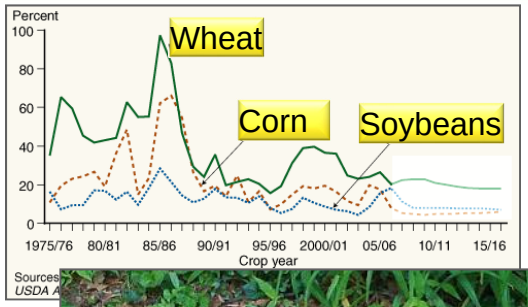

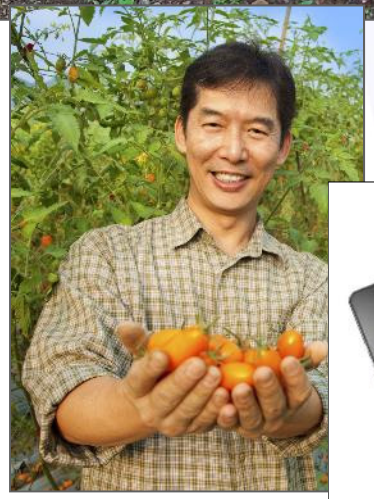

"A free world-class education for anyone, anywhere"

Dream #3: Affordable Cloud Solutions

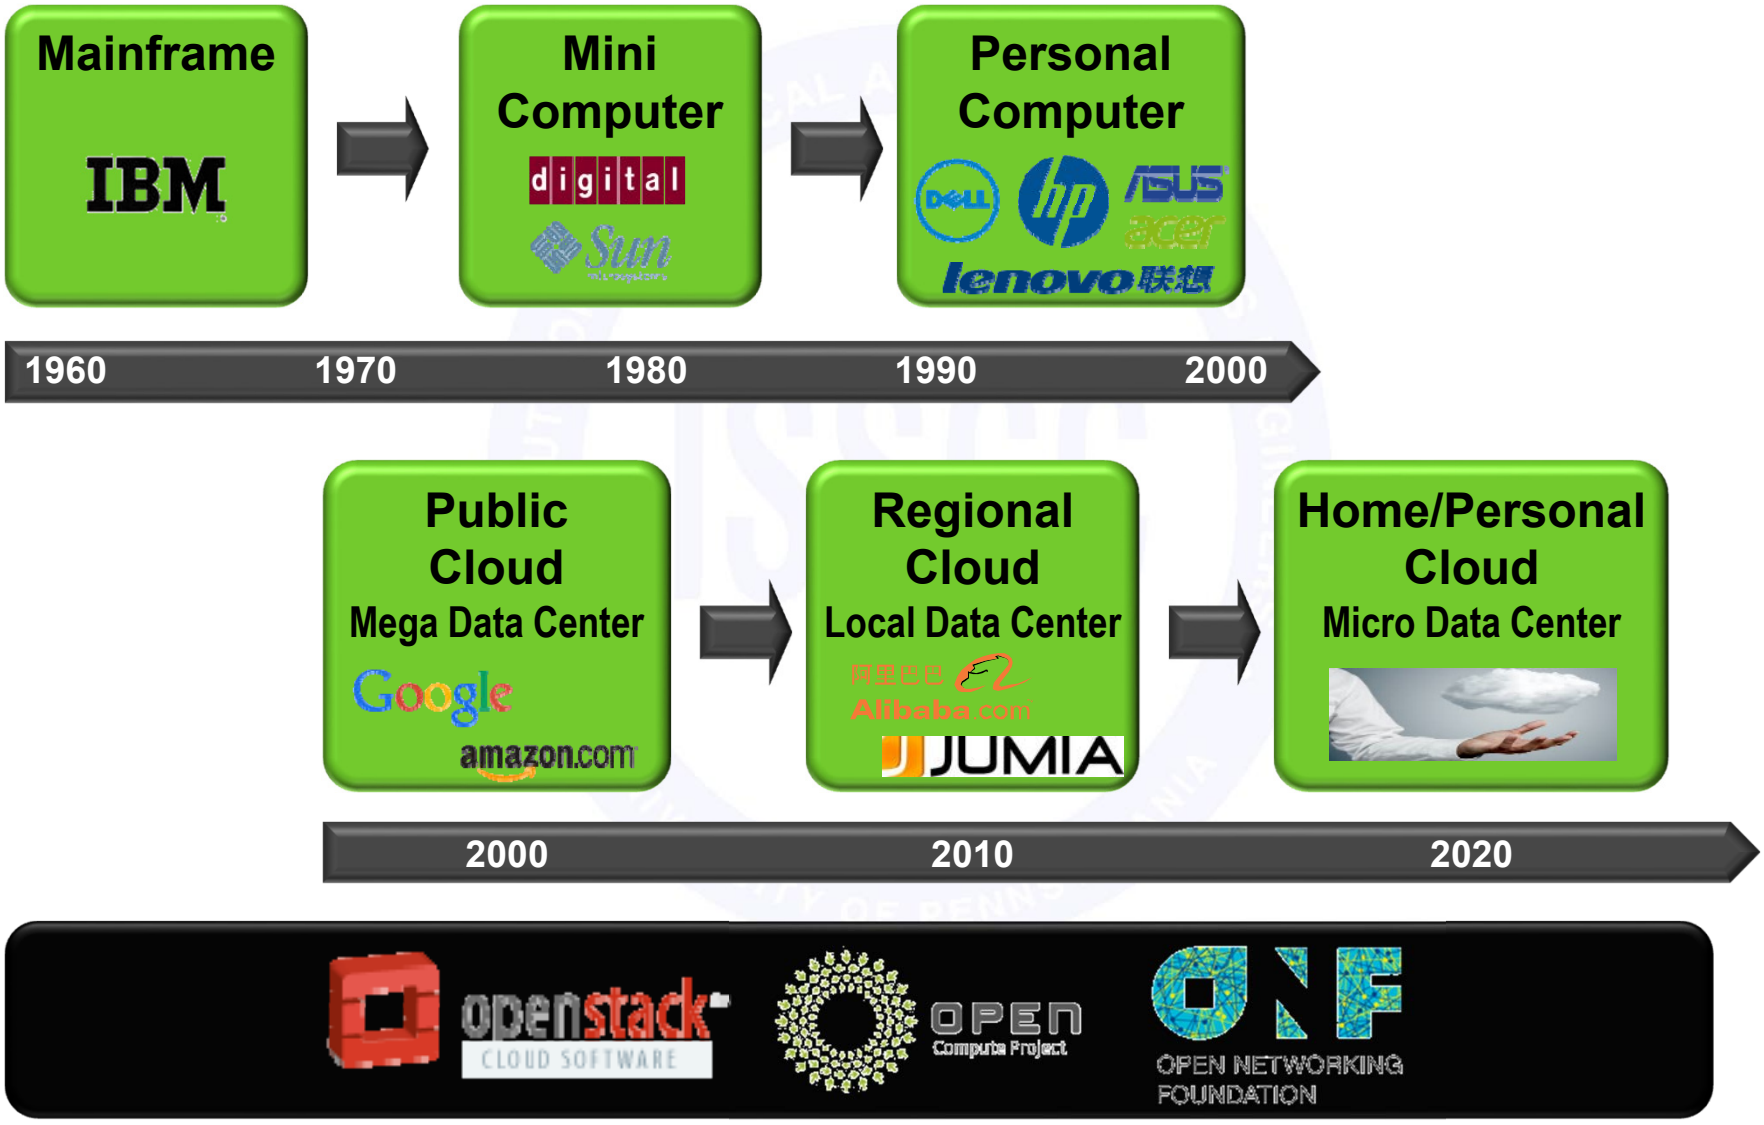

Next 5 Billion - Internet.org

Internet.org: Make internet access available to the other 5 billion people of the world

ERICSSON MEDIATEK OPERA software SAMSUNG

facebook NOKIA QUALCOMM

Cloud 2.0 Challenges (1/2)

ENERGY

→System/HW/SW
→Battery tech gap
→Energy harvest

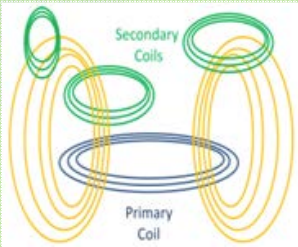

Wireless charging

Industrial solar

Heat-thermo

Vibration-piezo

BANDWIDTH

→Available spectrum
→Adaptive network
→Reliable connectivity

Micro

Small

Cloud 2.0 Challenges (2/2)

SECURITY

→Mobile health
→Mobile finance
→Analytics/big data

ELASTICITY

→Latency
→Compute capacity
→Workload

↔

Cloud 2.0: Ubiquitous Era

Cloud 1.0 triggers → Connected Devices… → Wearable & IoT, sensors → Local Clouds: Personal, home → Energy, Bandwidth Security, Elasticity

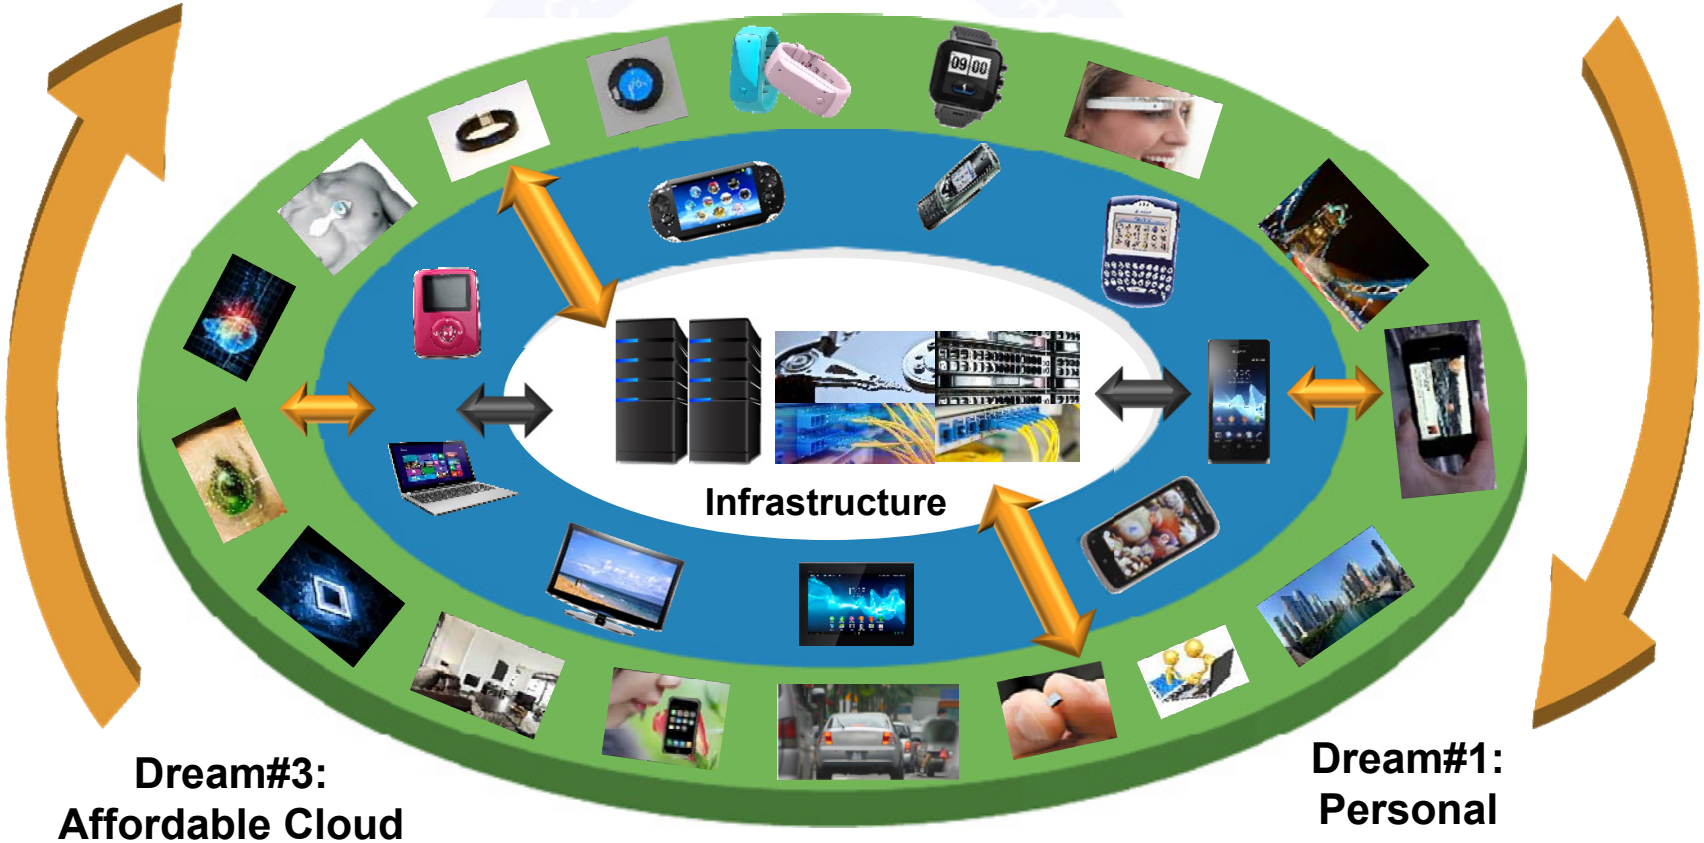

Dream#3: Affordable Cloud Solutions

Dream#2: Smart Devices for Everyone

Dream#1: Personal Assistant

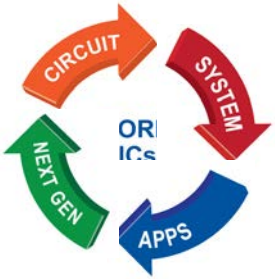

Conclusion

- **Driver:** Moore's Law & Positive Technology-Market Expansion Loop continue to drive computing cycles to Cloud 2.0 & beyond …It's upon us to further delay the "forever"
- **Cloud 1.0:** Innovations in performance and power made smart devices affordable
- **Cloud 2.0:** next wave of innovations are needed to drive 10X lower cost of cloud services everywhere
- **Impact:** Cloud 2.0 will potentially have as big an impact in changing our life style as auto industry had in the 20th century

"Enhance and Enrich Everyone's Life"

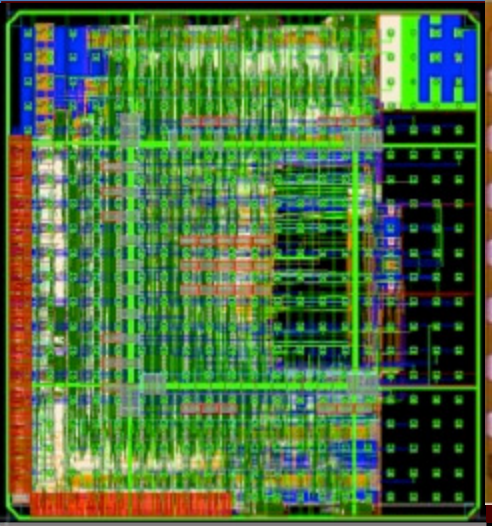

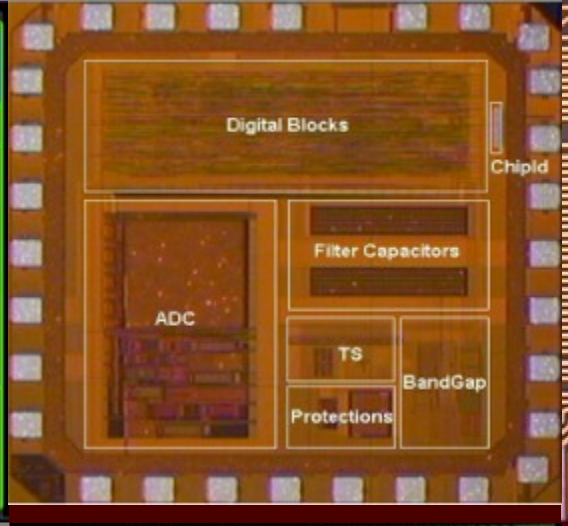

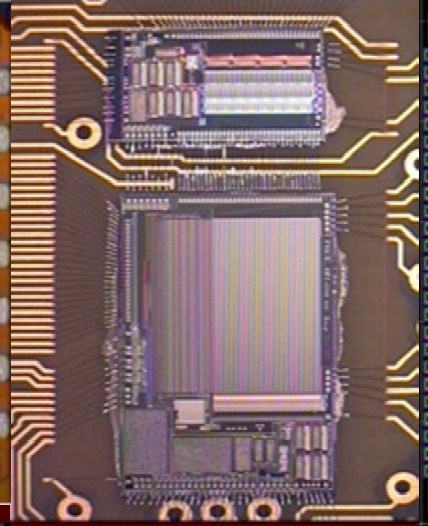

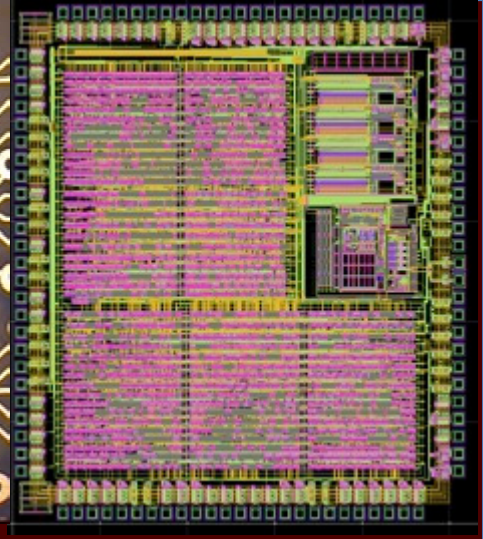

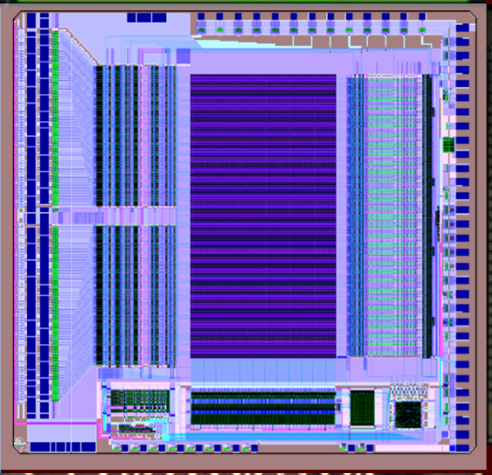

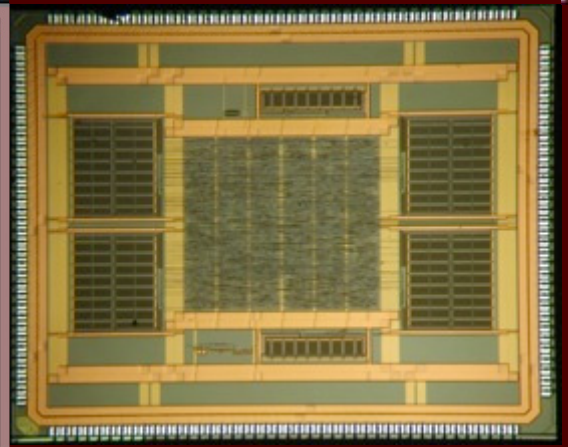

How Chips Pave the Road to the Higgs Particle and the Attoworld Beyond

Erik H.M. Heijne

CERN – PH Dept. CH 1211
Geneva 23

Institute for Experimental and Applied Physics of the Czech Technical University in Prague

Nikhef Amsterdam

Erik.Heijne@cern.ch

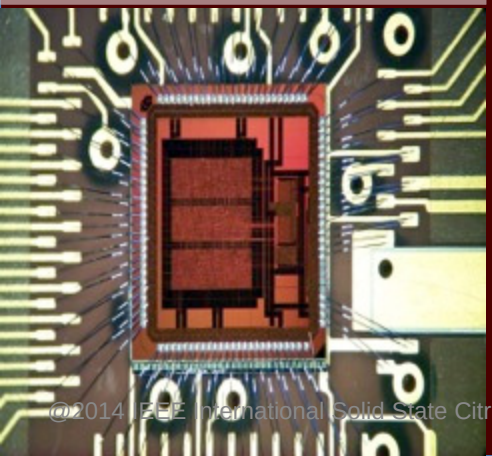

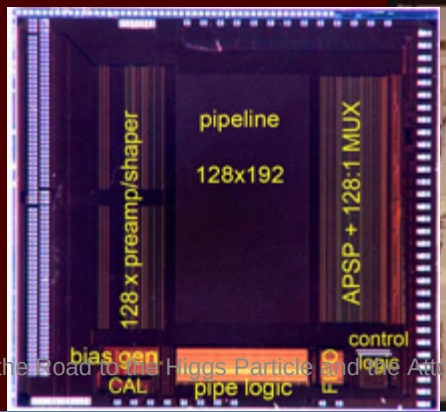

Acknowledgements

Thanks to:

CERN Microelectronics Group Philippe Farthouat, Alessandro Marchioro,
Michael Campbell, Pierre Jarron, ...

Leuven - ESAT and IMEC Switzerland EPFL - ETHZ

Many specialists around the world

A. Zichichi and LAA project helped start chip design at CERN

Governments who support Science
at CERN and in their countries

ISSCC reviewers

Periodic Table based on “Standard Model” Basic Forces and Constituents of Matter

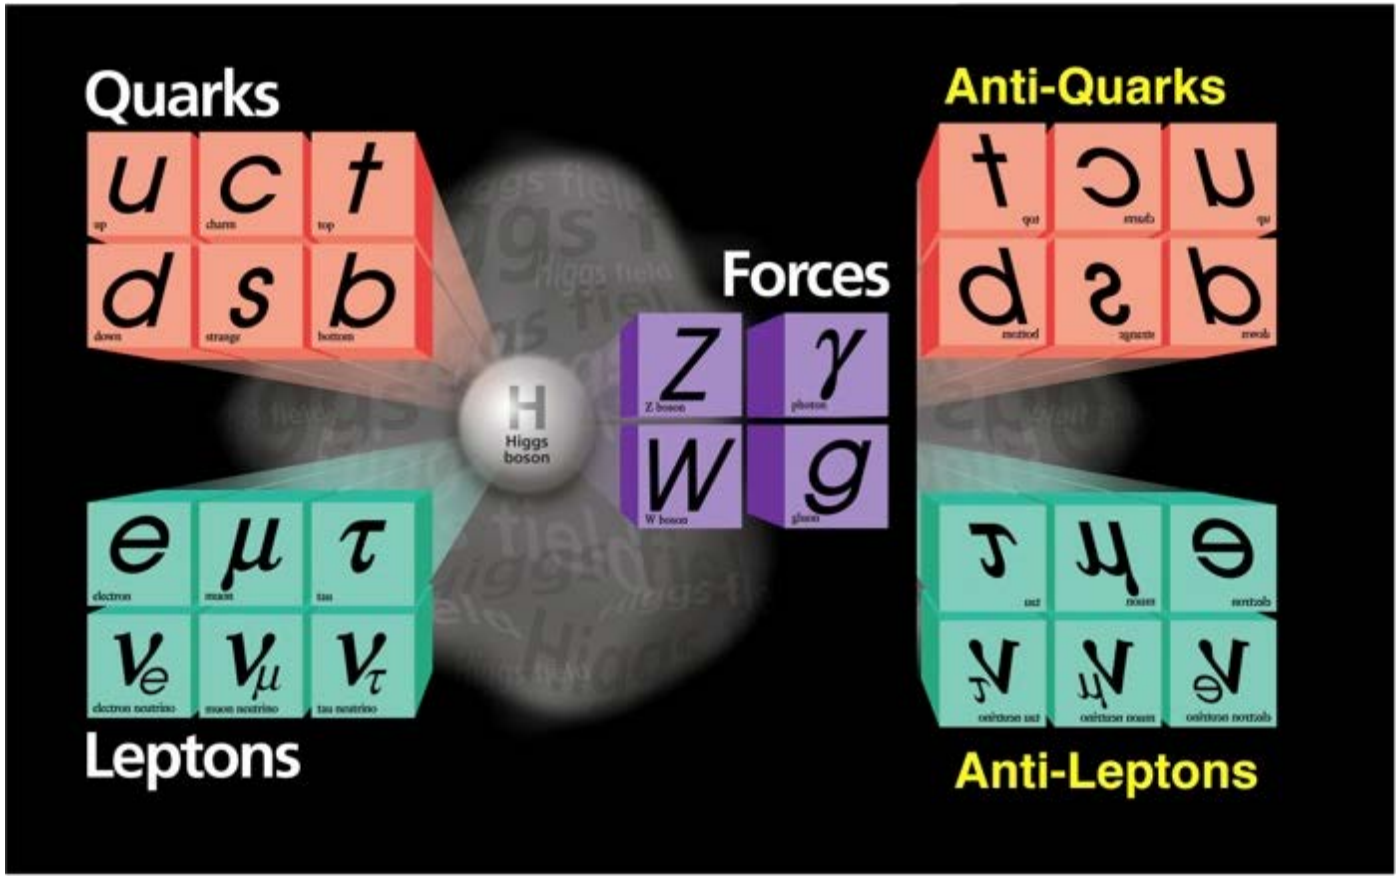

source: Fermilab

Many questions remain:

- why 3 generations
- where are the original antiparticles
- how to include gravity
- .....

The mass problem

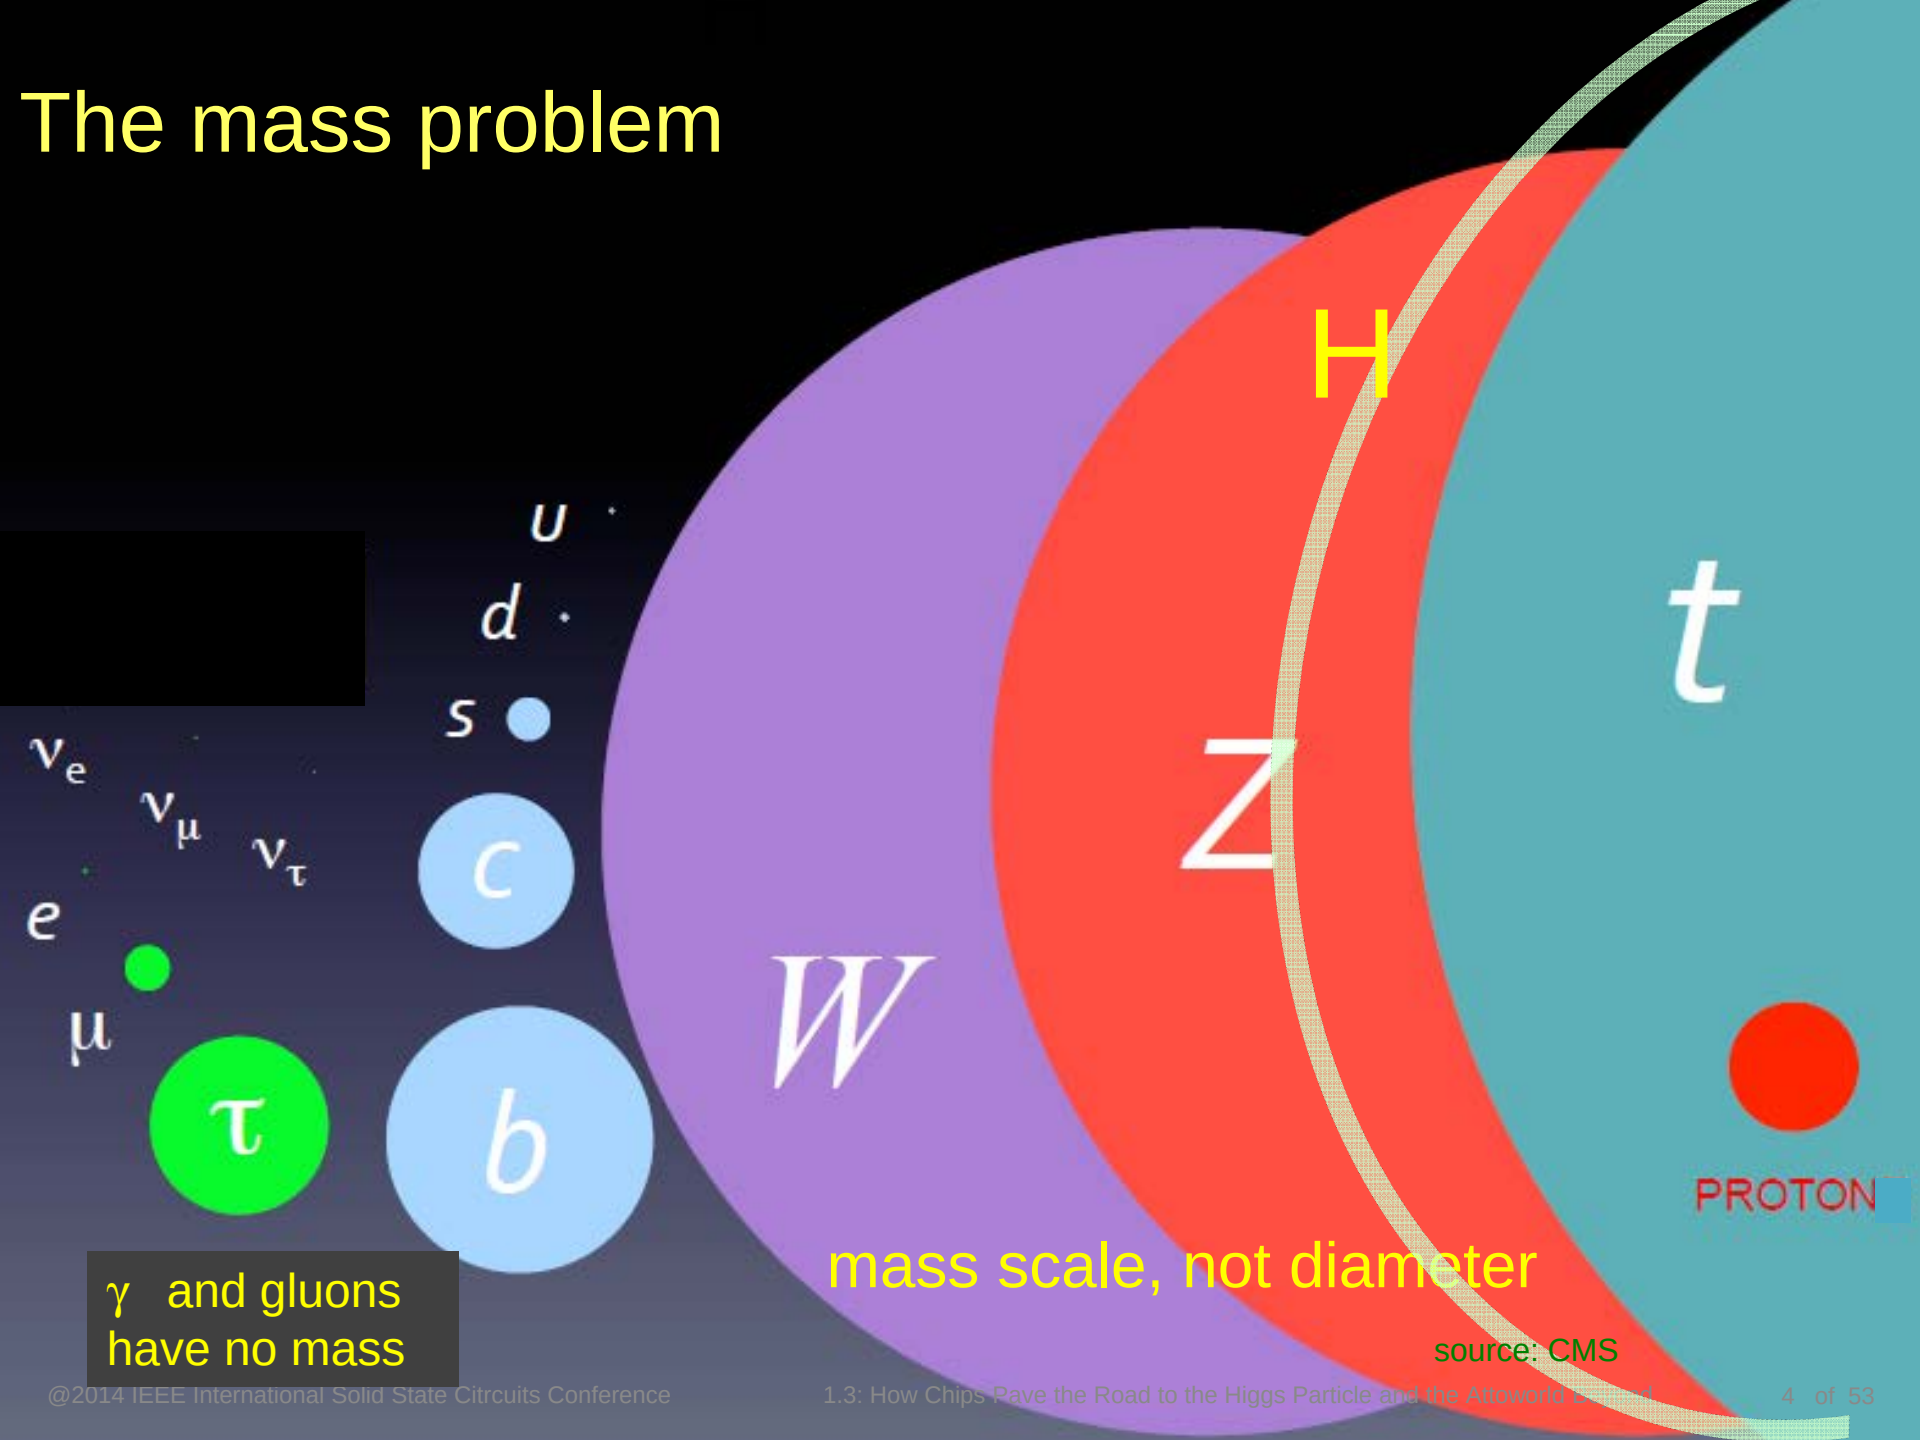

Energy, Mass and Wavelength of Particles

Particles are Quanta of Energy

$E=mc^2$

At the same time described as

Waves with Frequency in a Field

$E= \hbar v$

TeV energy → attometer wavelength

nanoelectronics 10^{-9} m

attoworld 10^{-18} m

many atto particles exist only briefly

Shorter Wavelengths, Larger Imagers

microscope
visible light
photons ~1eV
size 30 cm
µm objects

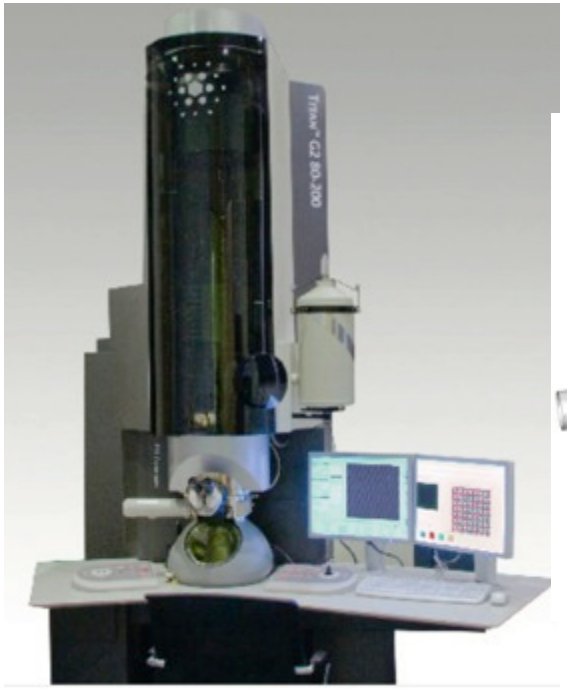

TEM transmission
electron microscope
electron beam ~100 keV
size 3 m
Å objects

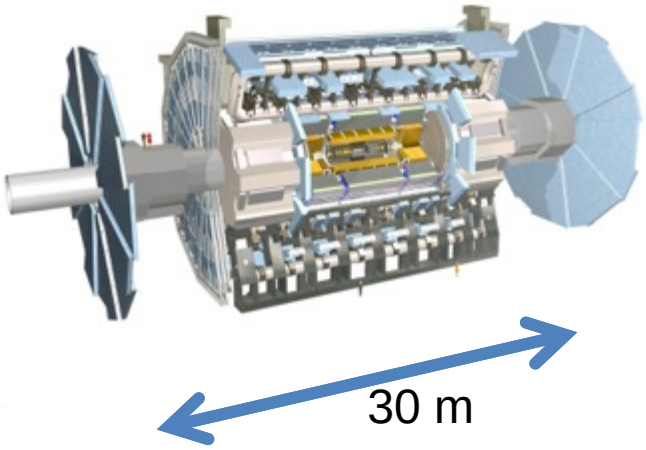

'attoscope' CERN
using protons 7 TeV

30 m + 30 km ring
atto meter objects

2008 Large Hadron Collider LHC at CERN Geneva

In tunnel ~100m underground

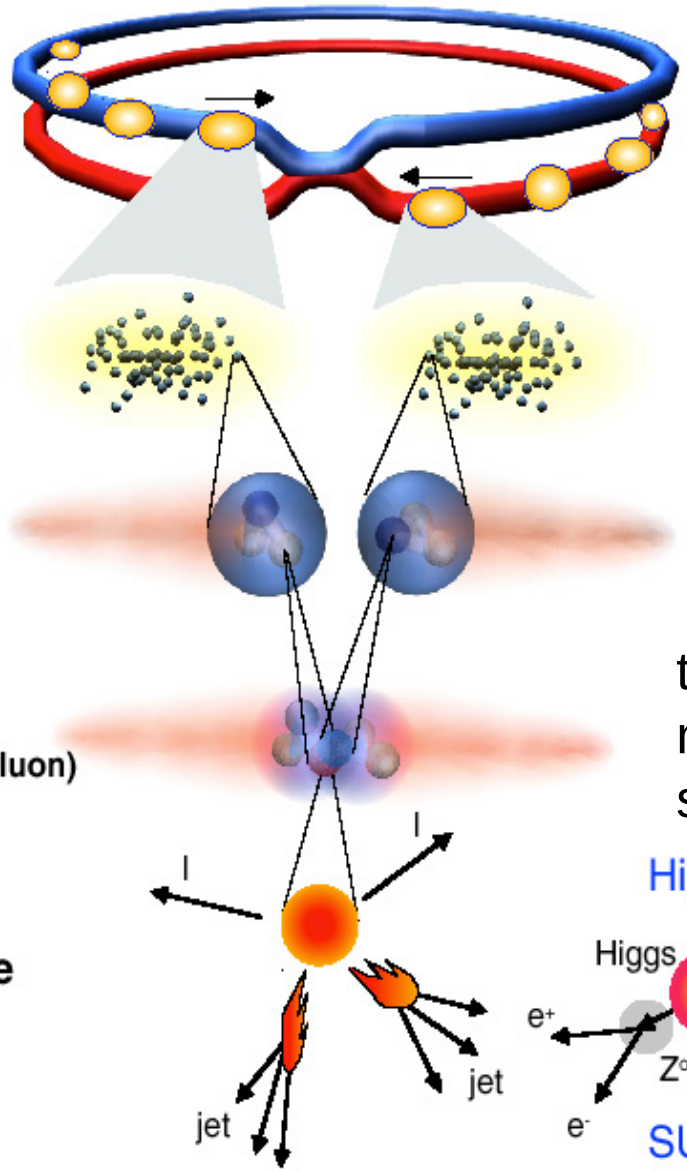

Proton-Proton	2835 bunch/beam
Protons/bunch	10^{11}
Beam energy	7 TeV (7×10^{12} eV)
Luminosity	10^{34} cm^{-2} s^{-1}

Crossing rate	40 MHz

Collisions ≈	10^{7} - 10^{9} Hz

the heavy quanta
must be created in situ
so that they can be studied

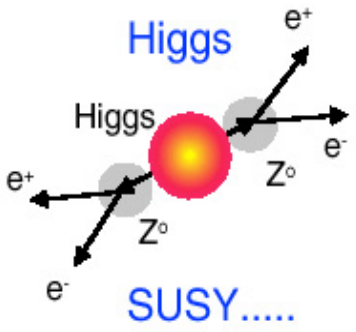

Selection of 1 in 10^{-13}

10,000,000,000,000

~1970 Discover Particles by Imaging in Bubble Chamber

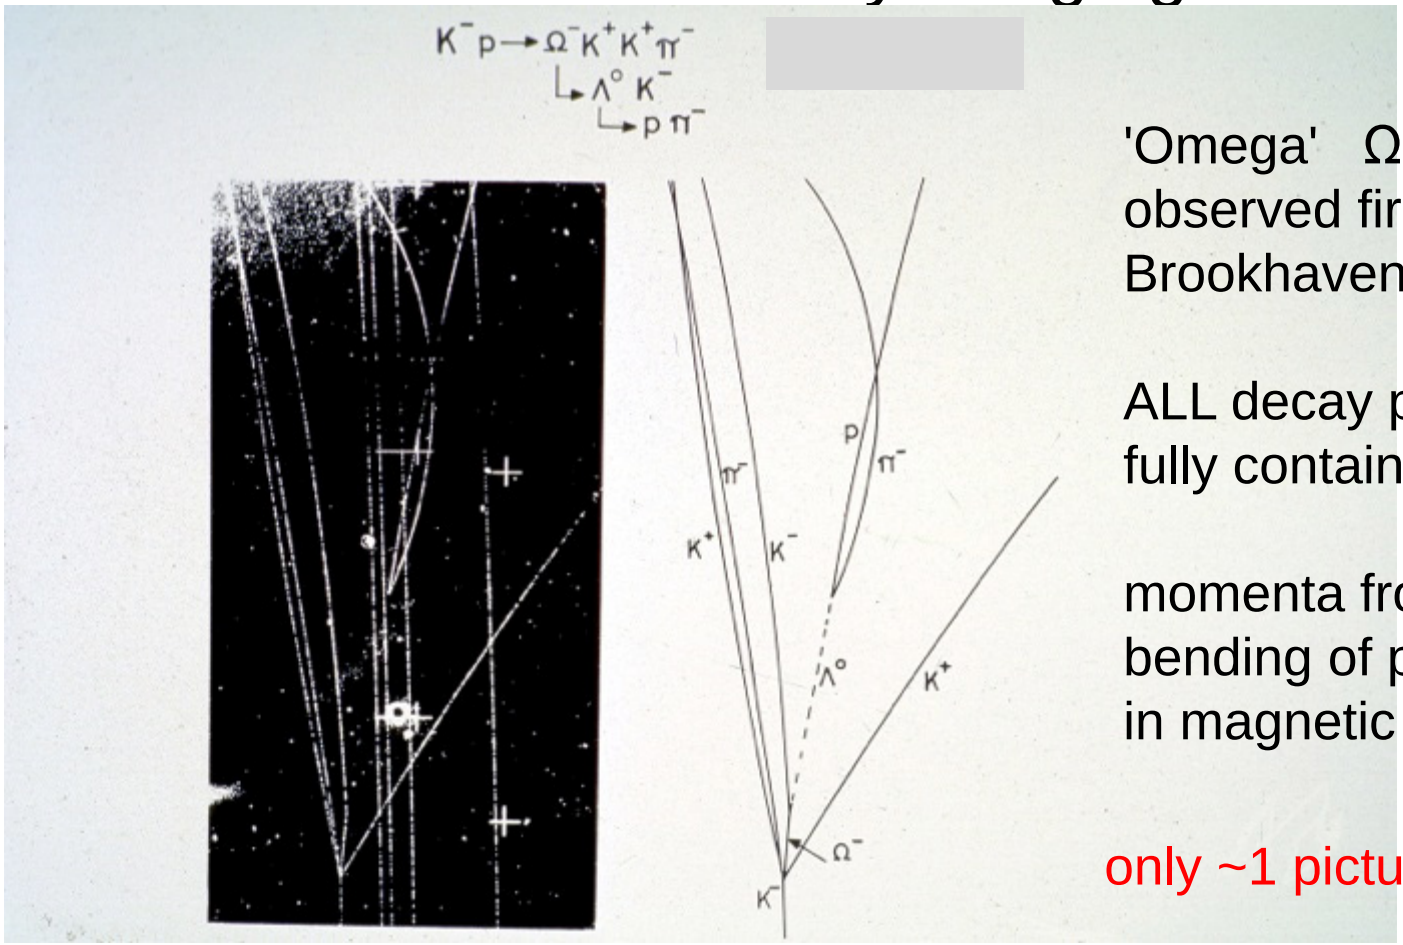

'Omega' Ω particle
observed first in 1964
Brookhaven Nat Lab

ALL decay products
fully contained

momenta from
bending of particles
in magnetic field

only ~1 picture per second

3 photo-views in CERN
2-m bubble chamber

10 GeV
incident protons

Imaging Now
All Electronic
with 3-D
Reconstruction

Many Tracks
and 2 “Jets”

40 million / sec

Secondary
Vertex:
a short-lifetime
particle is a
messenger for
something new

see blow-up

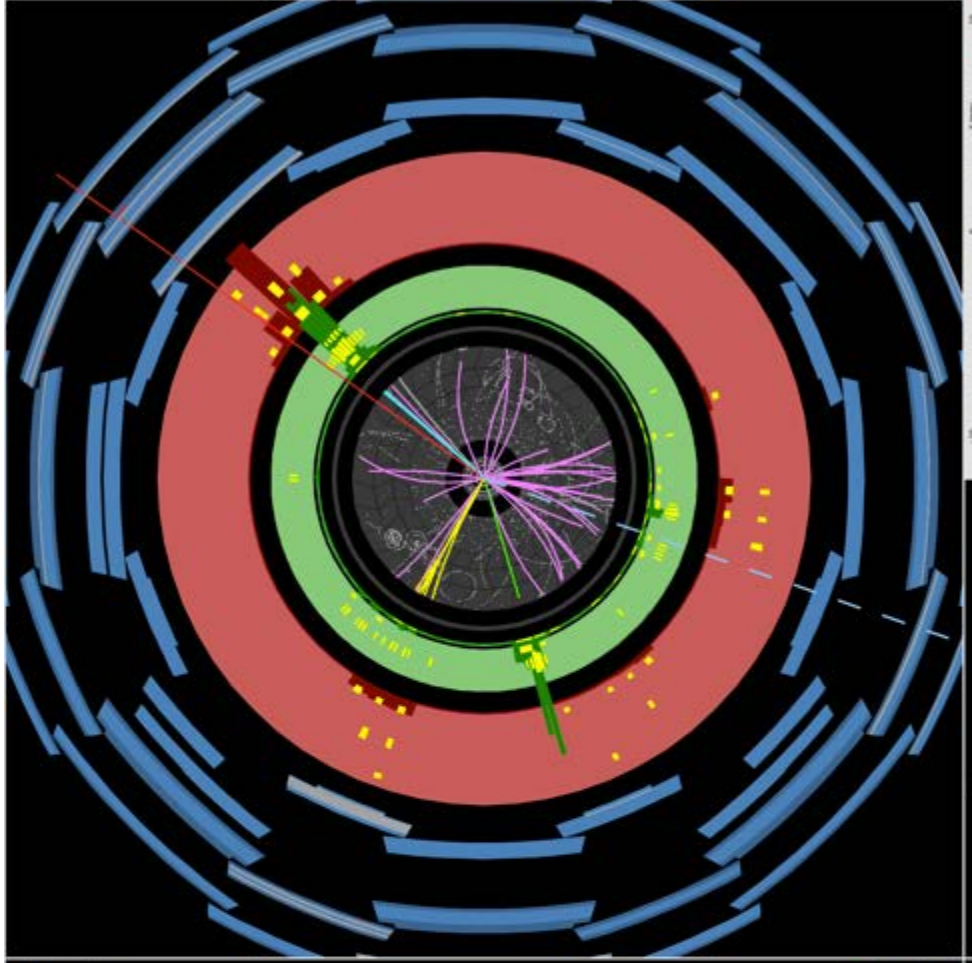

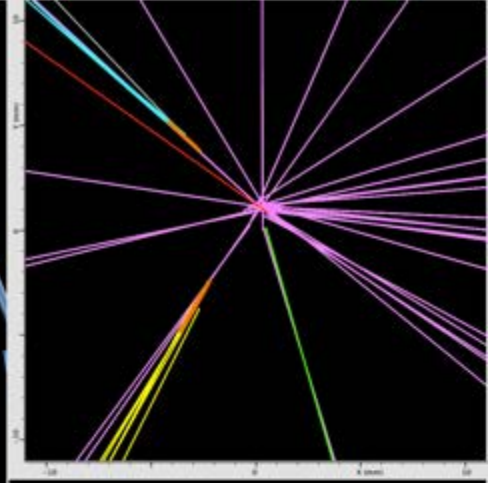

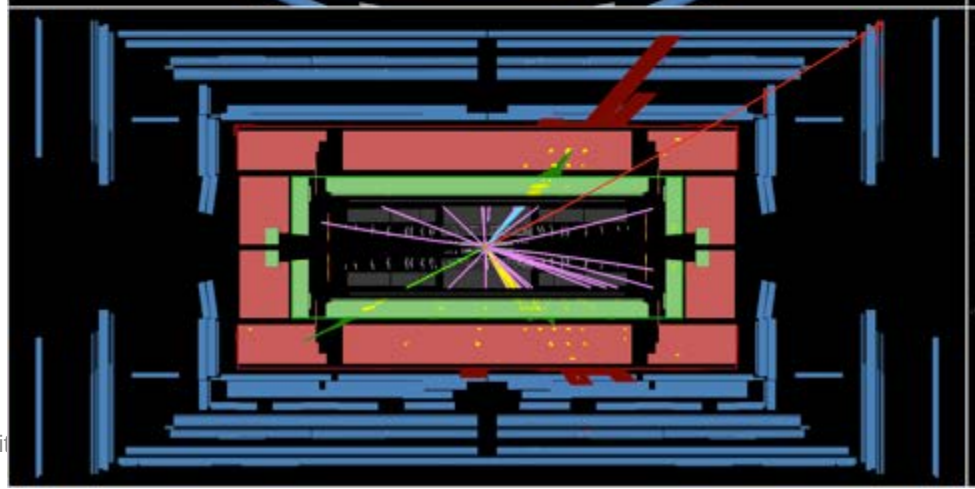

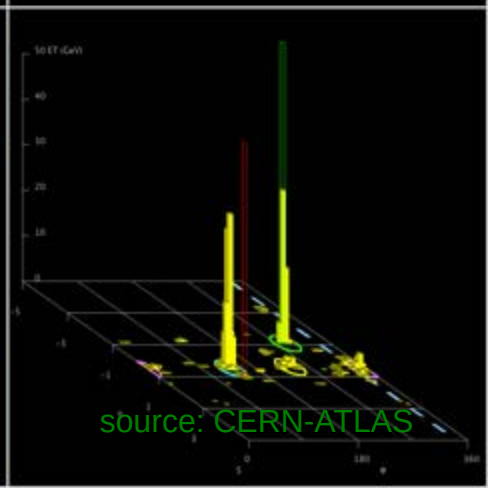

ATLAS

Details around Primary Vertex

Two Secondary Vertices

Note scale
1cm
all this is INSIDE beam pipe ∅ 7cm

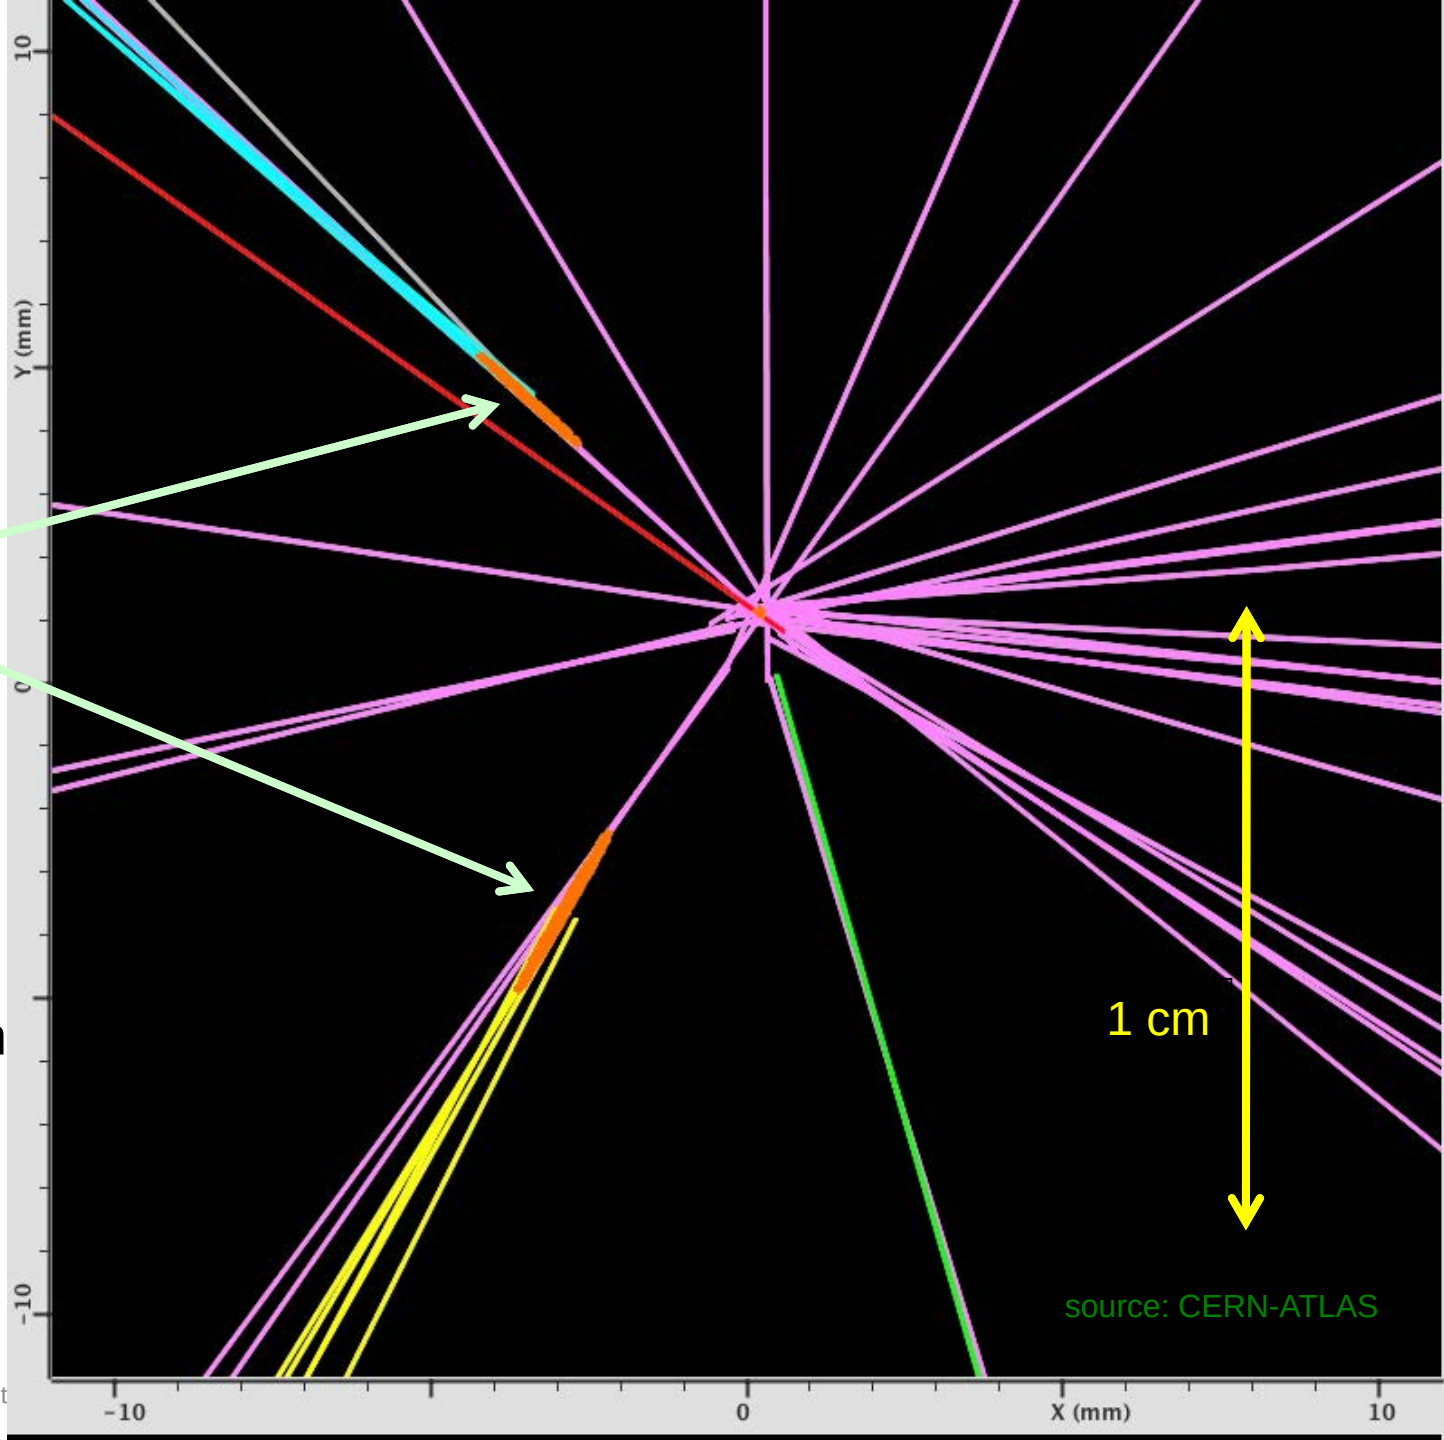

Probability for Interactions

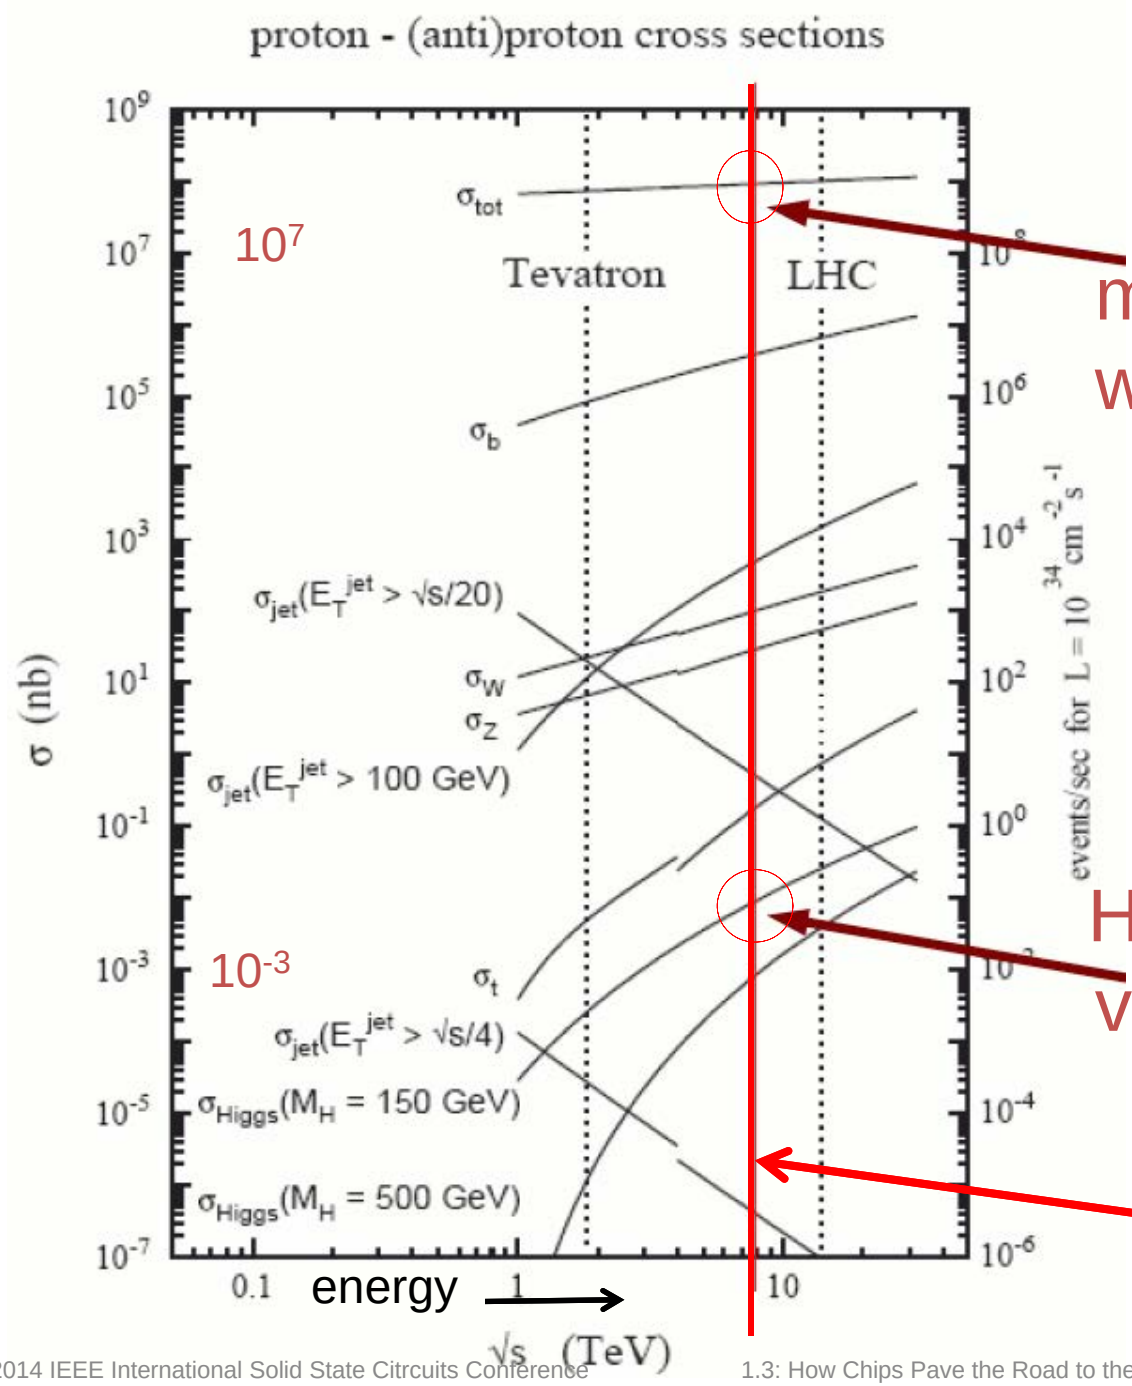

mostly produce well-known objects

Higgs production very rare: $<10^{-10}$

LHC energy 4 + 4 TeV

Schematic Data Processing and Filtering

3 main sensor layers

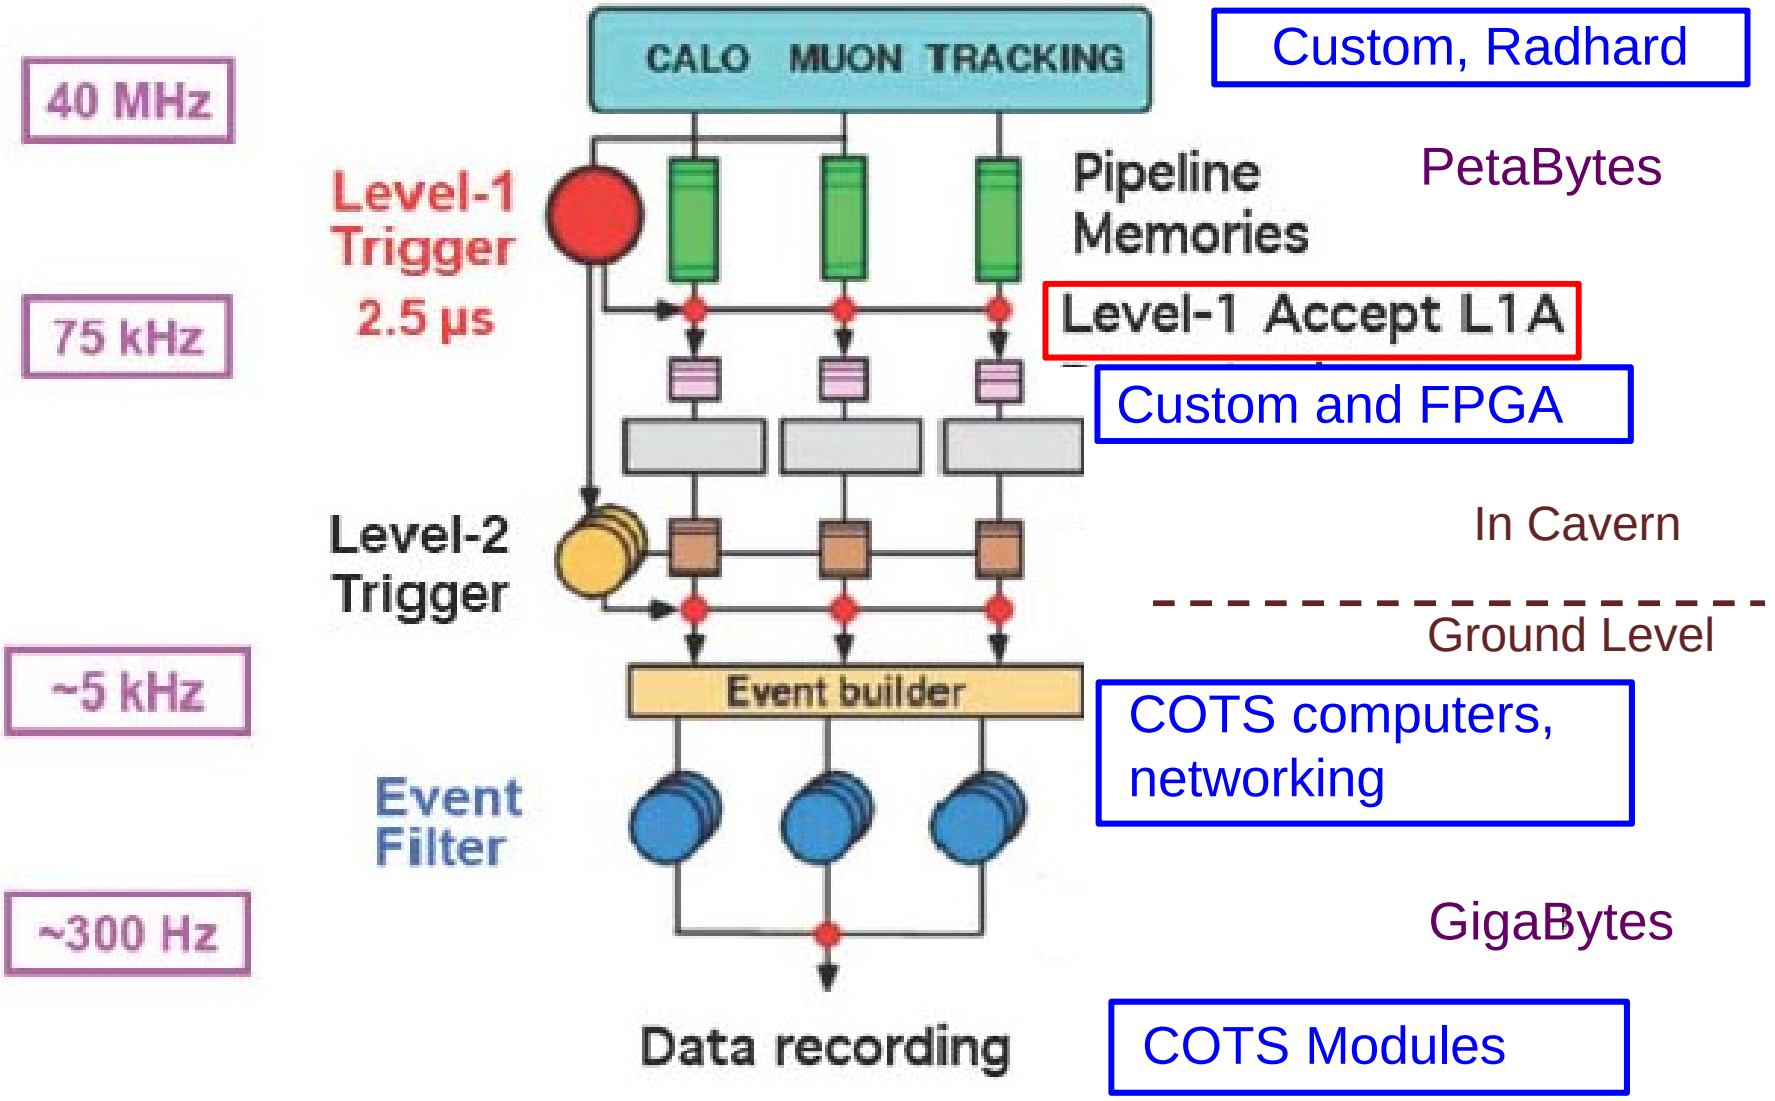


source: CERN-ATLAS

Animation of CERN Accelerators

screenshot

68 sec

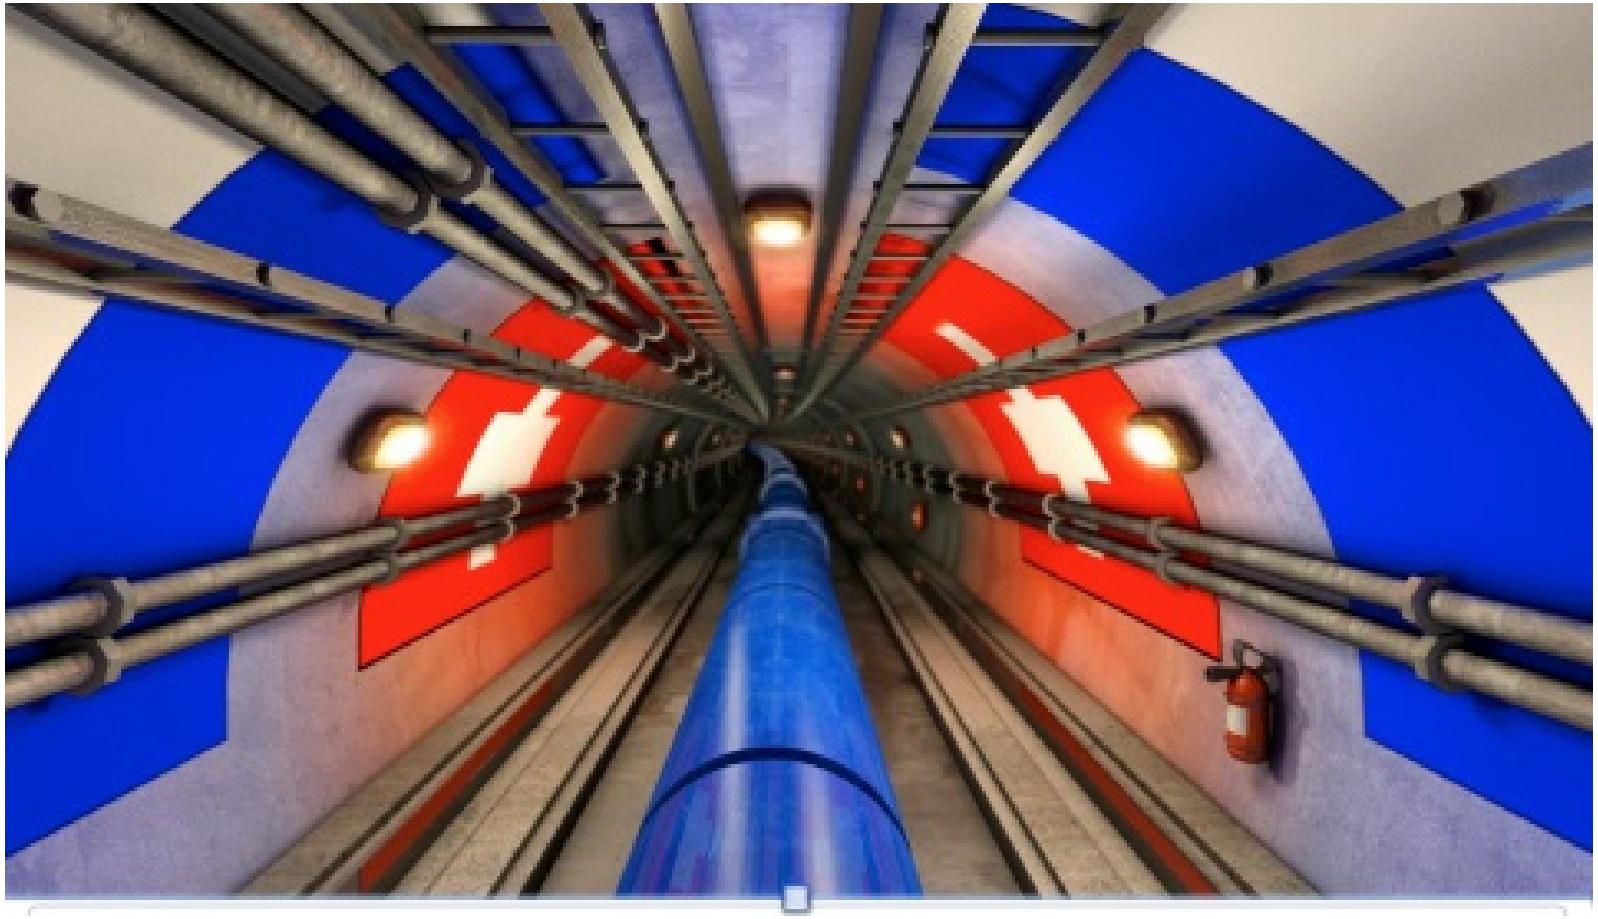

Experiment is Giant 'Camera' with Magnetic Field

Experiment is Giant 'Camera' with Magnetic Field

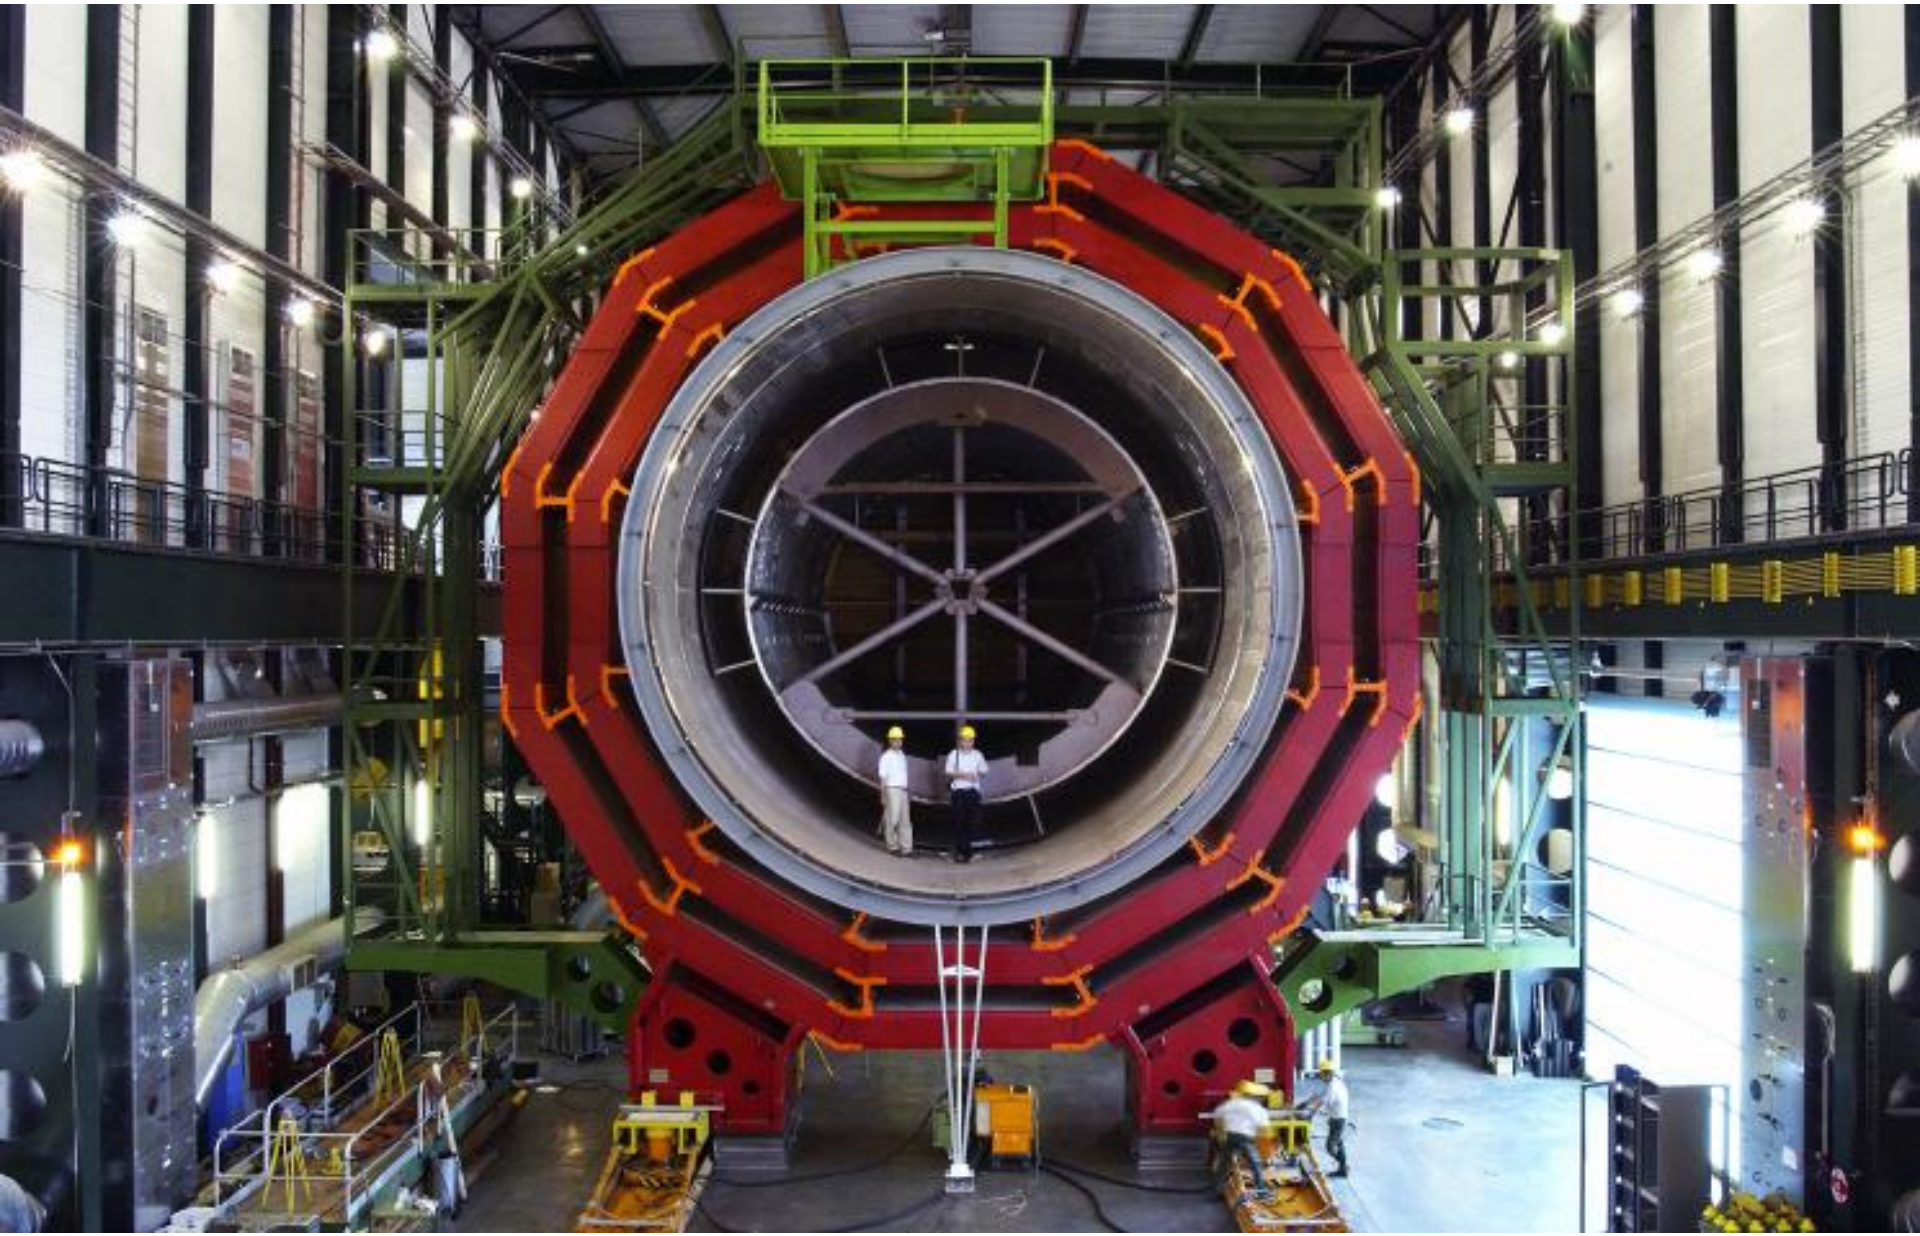

CMS Collaboration

Two Types of Si Sensors

Si Microstrip Detectors
can cover larger area

Linear array
diodes 2-15 cm
depleted bulk ~300µm

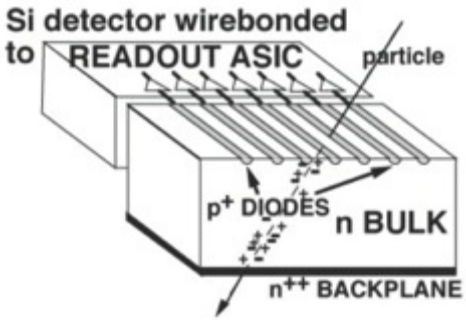

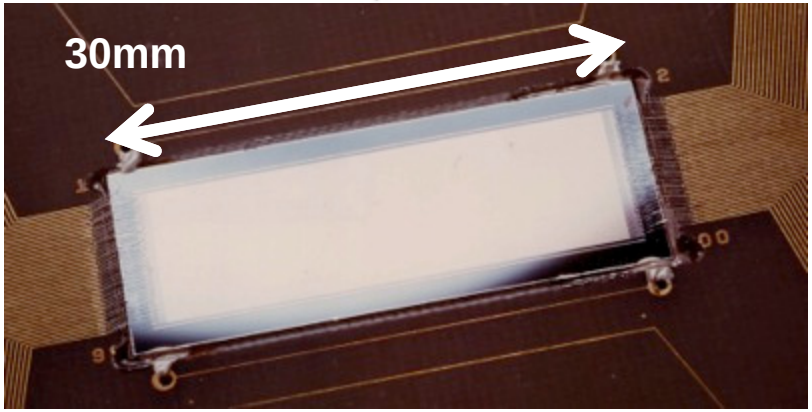

Si Pixel Detectors
for highest densities

2D pixel matrix with bumps
diodes 30µm – 500µm
depleted bulk ~150µm

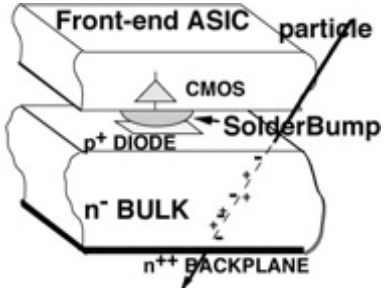

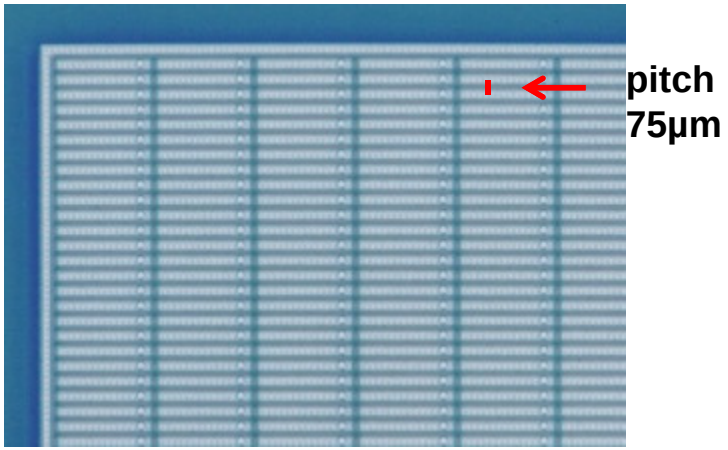

Why Silicon Sensors Save the Day

Signal ~70 e^- per µm →20 000 e-h pairs
proportional to particle energy loss

Signal consist of electrons and holes

Both are collected with ~comparable speed

Sensitive volume is clean again after ~15ns

Lithography gives ~µm precision on track

Detector is Built-up in Layers : e.g. ATLAS

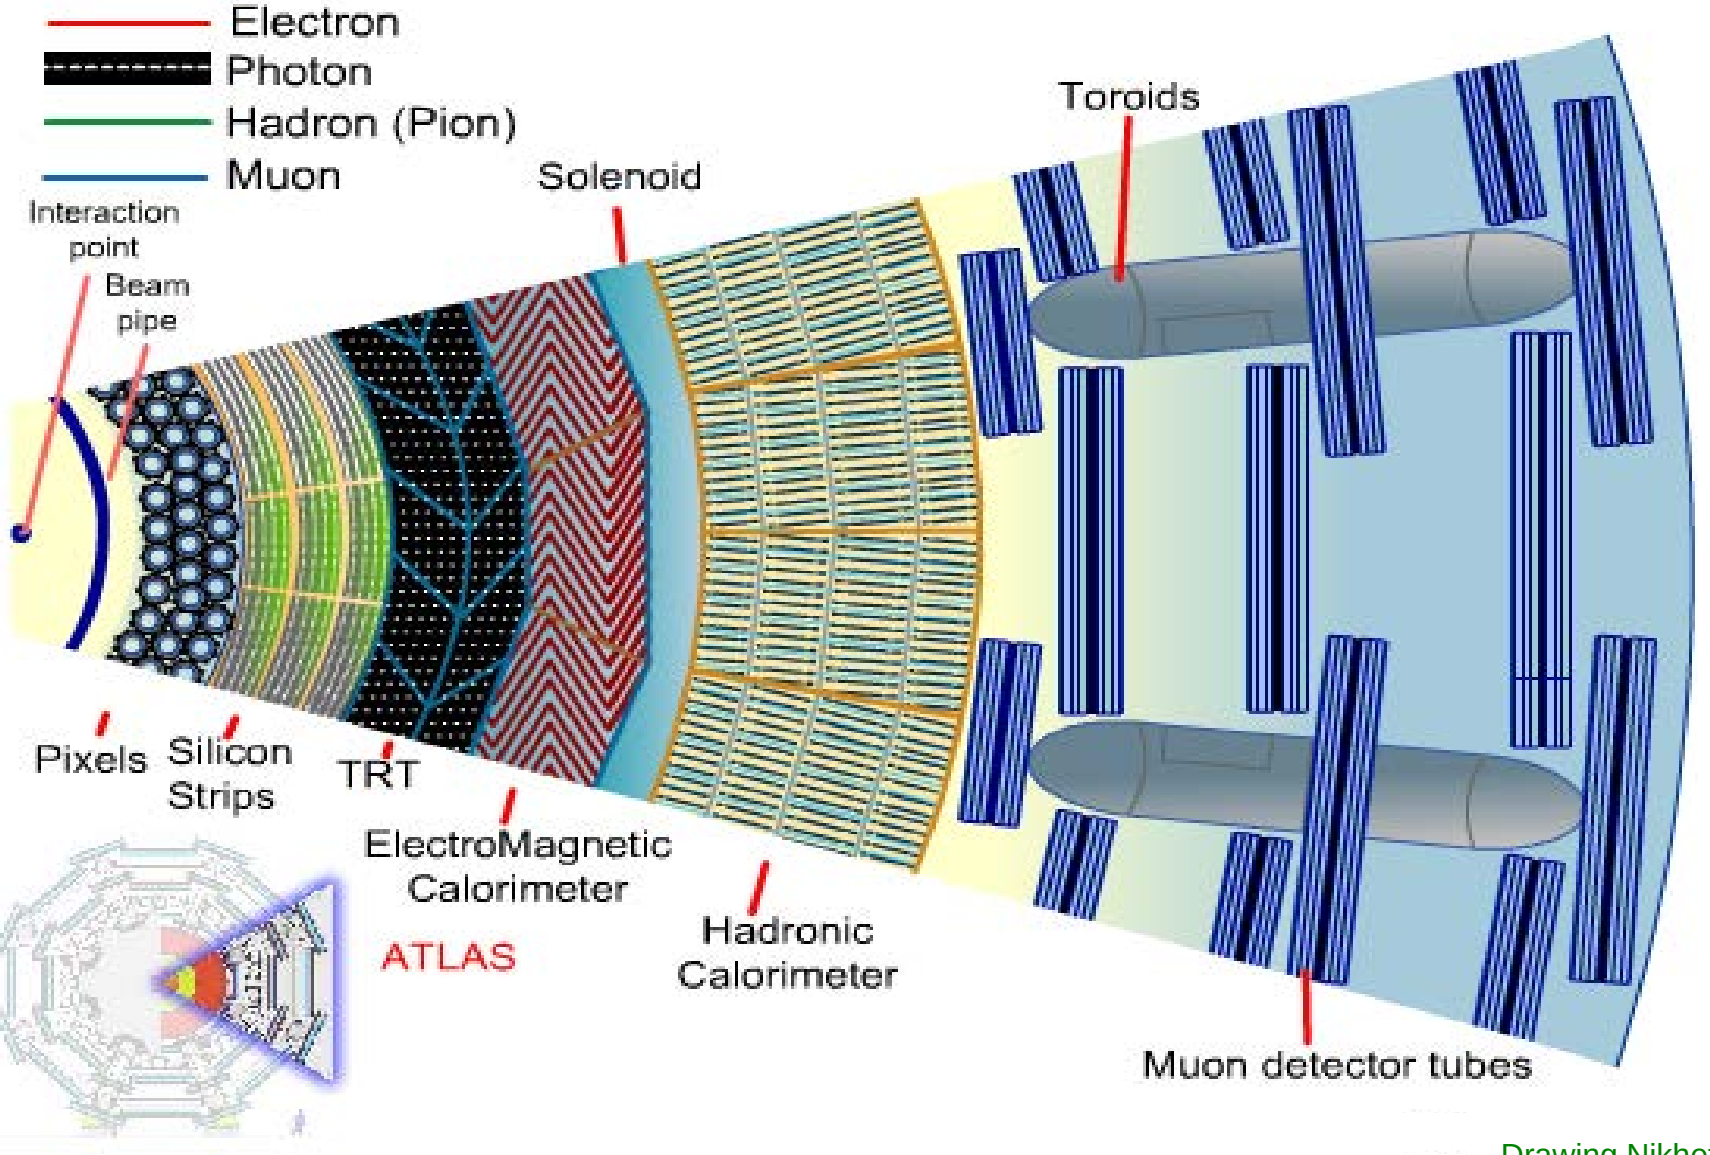

Drawing Nikhef

Animation of function of detector layers

screenshot

53 sec

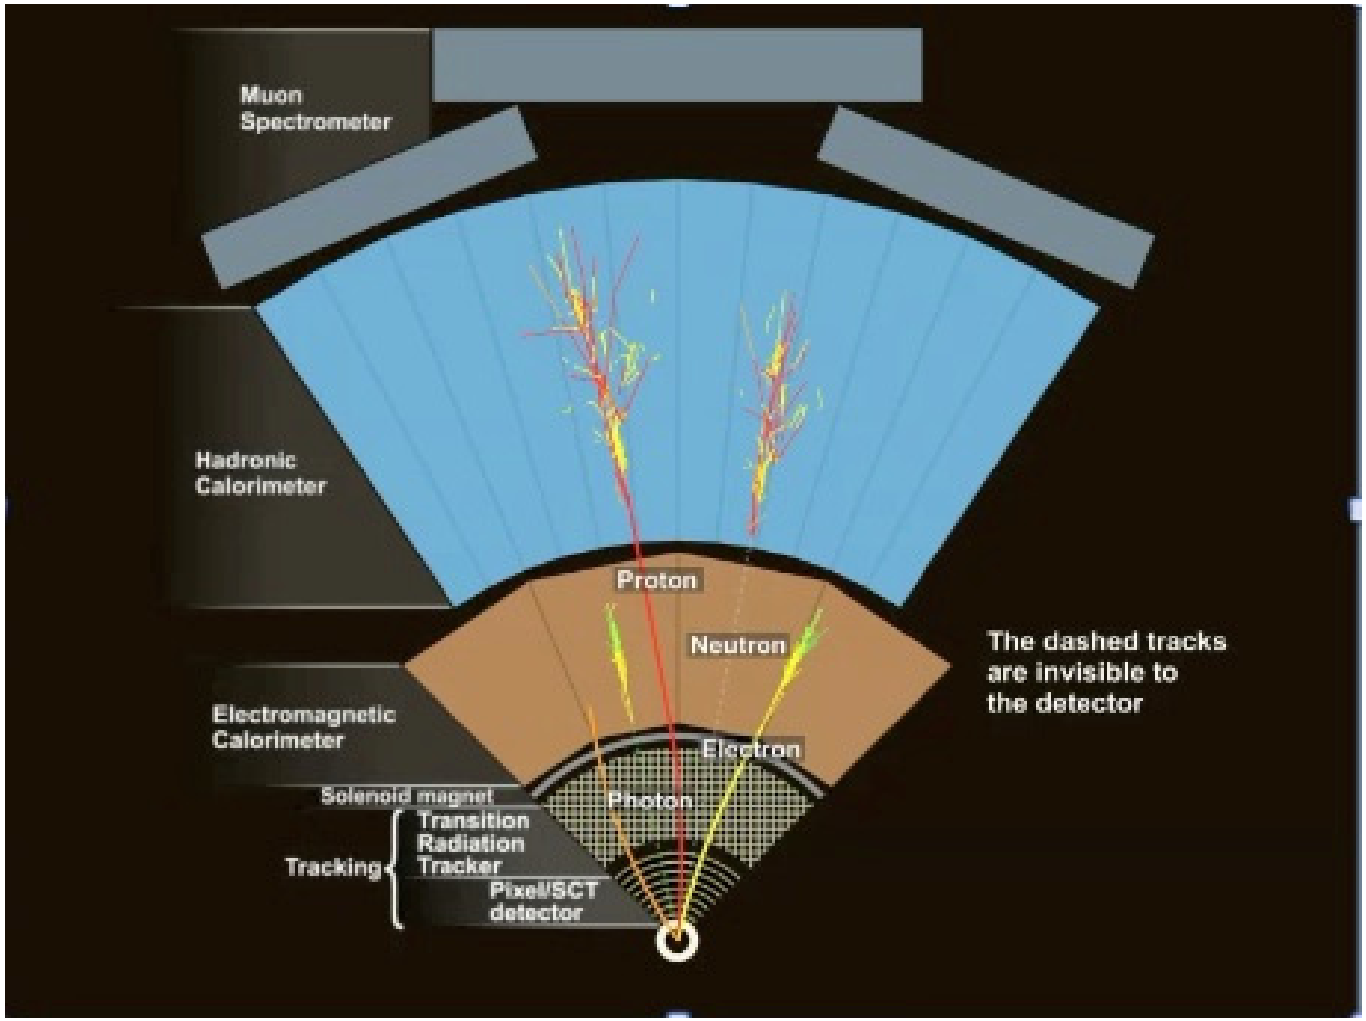

CMS Silicon Tracker

HALF- BARREL

Inner barrel half-module under construction

Si-sensors and chips

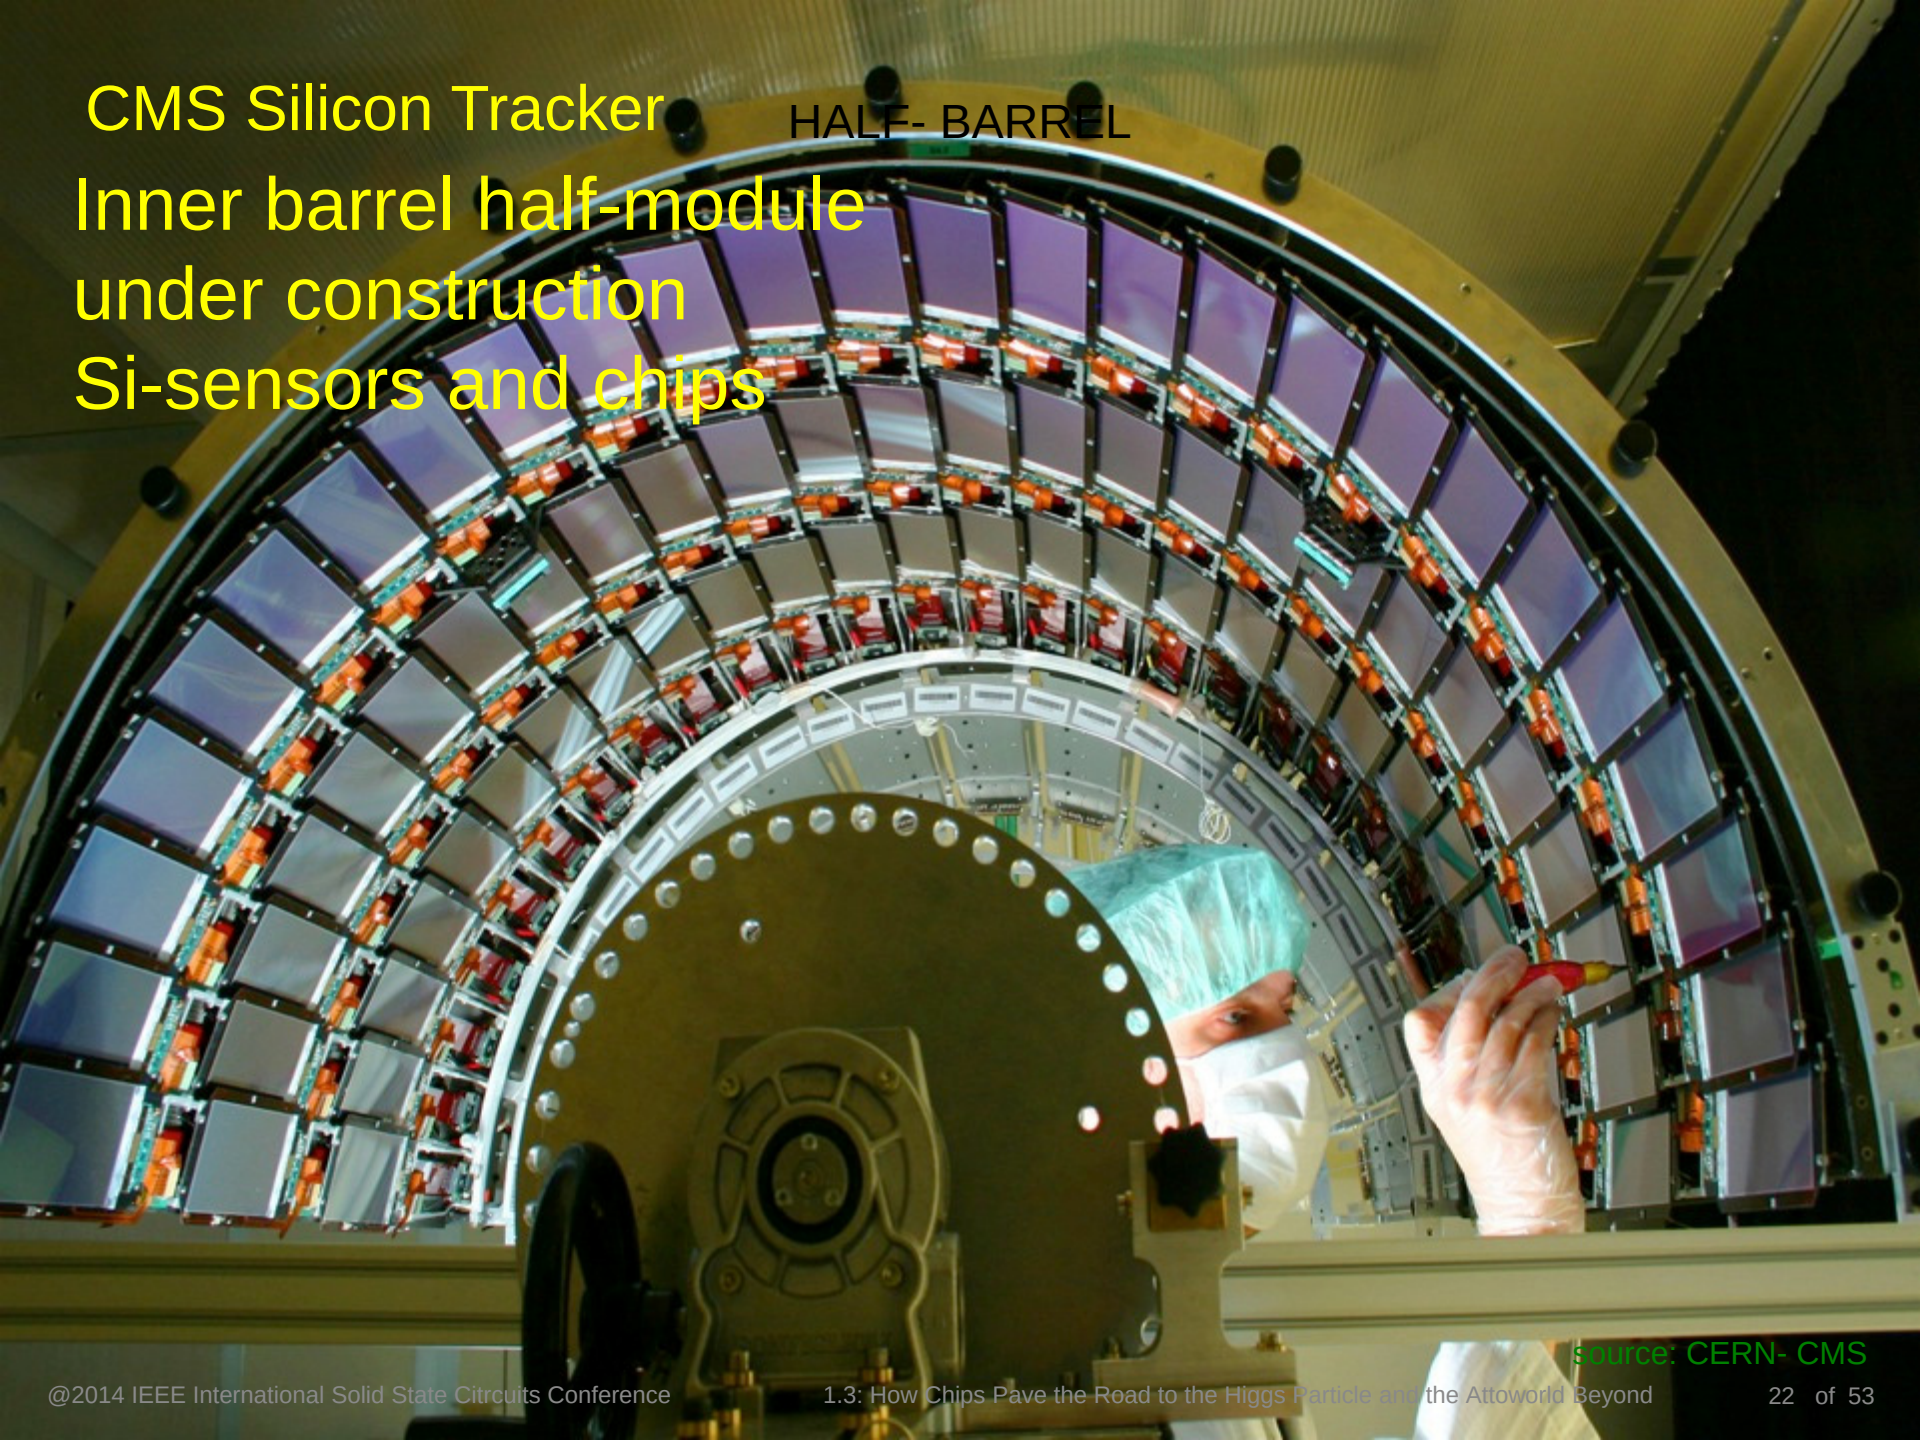

source: CERN- CMS

source: CERN- CMS

Insertion of CMS Si Tracking Detector

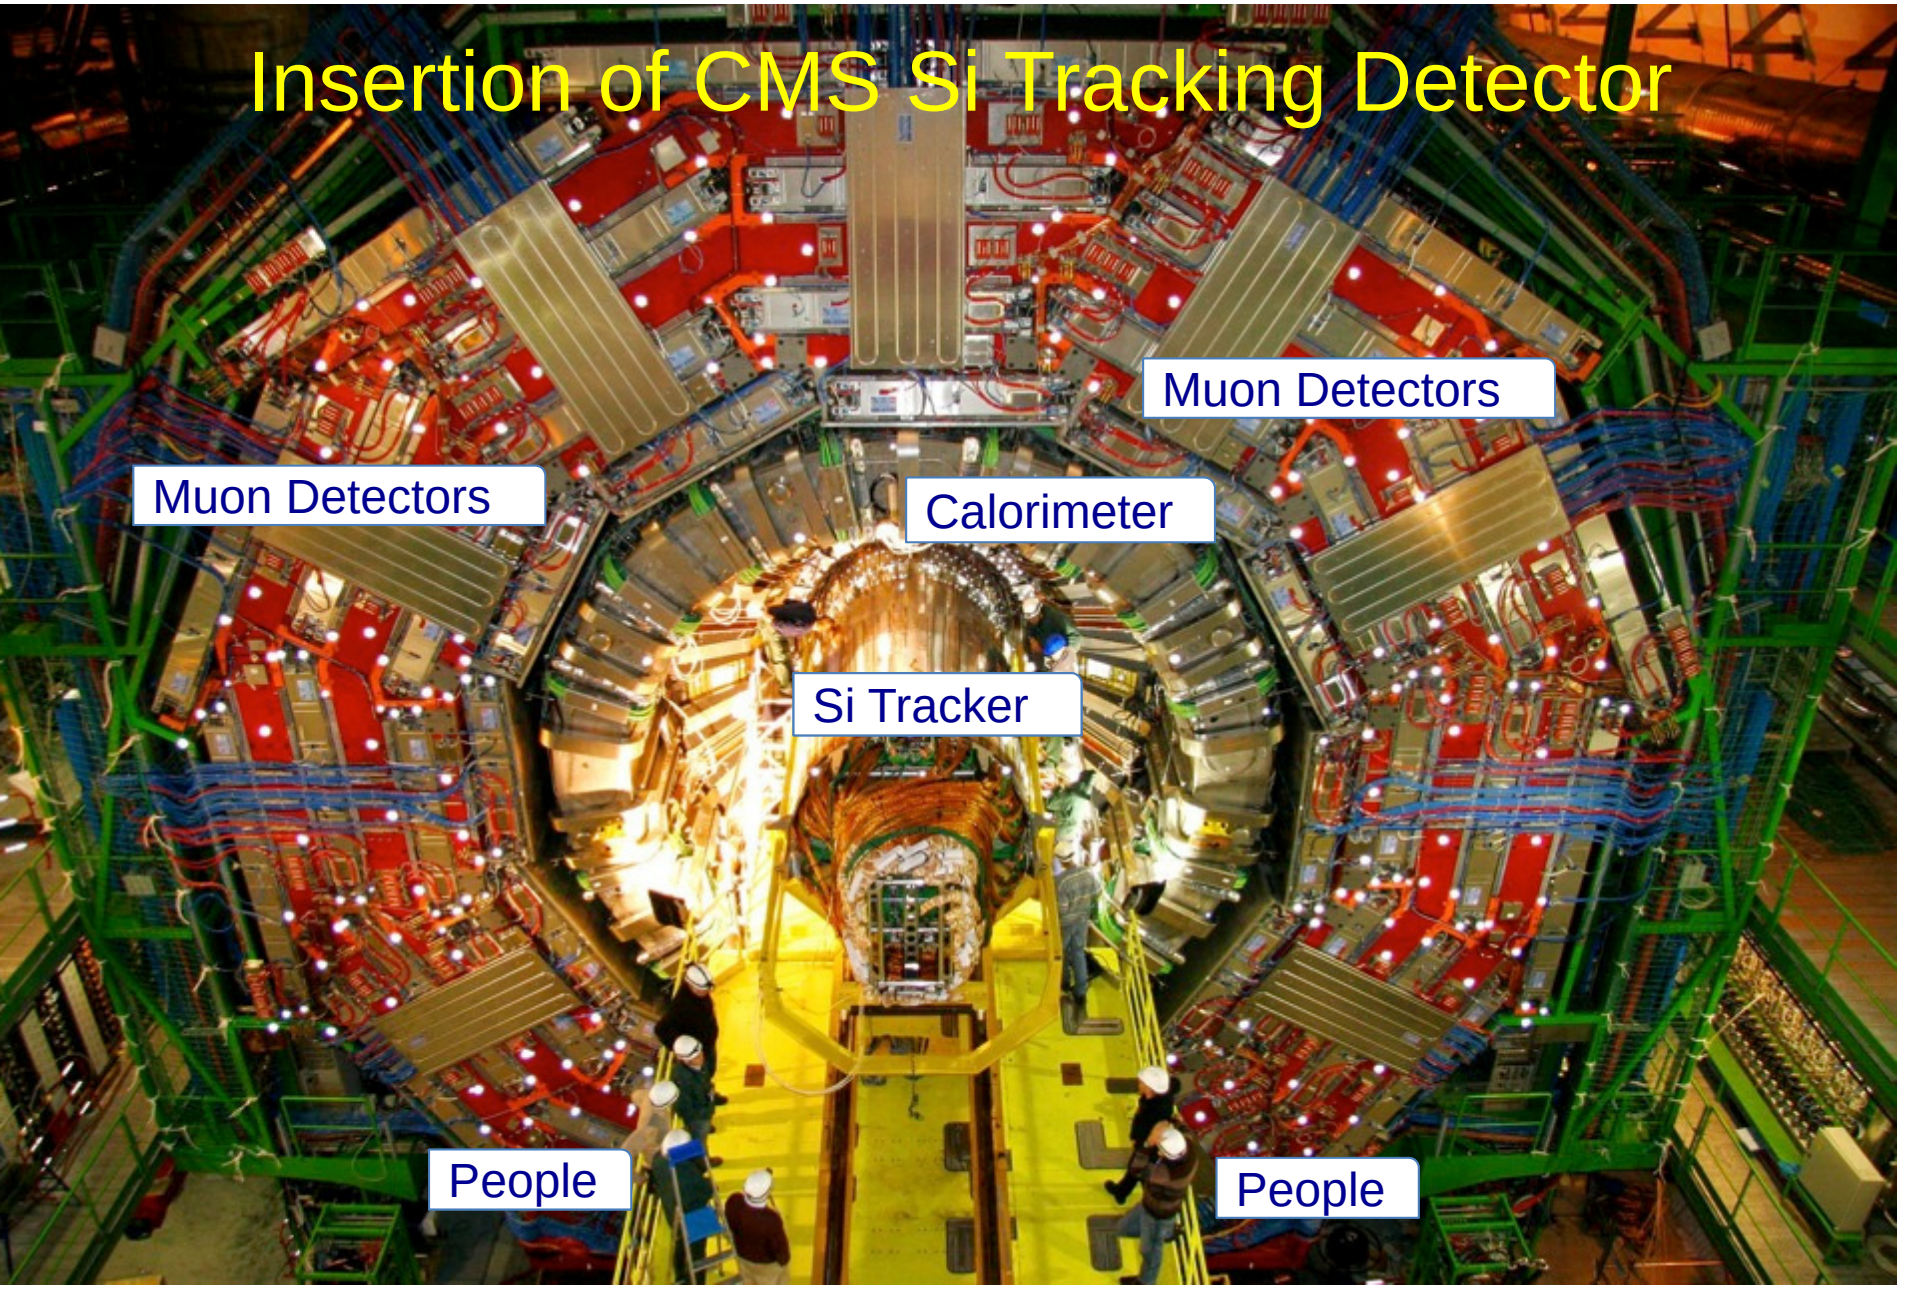

source: CERN- CMS

Some of the ASICs in ATLAS

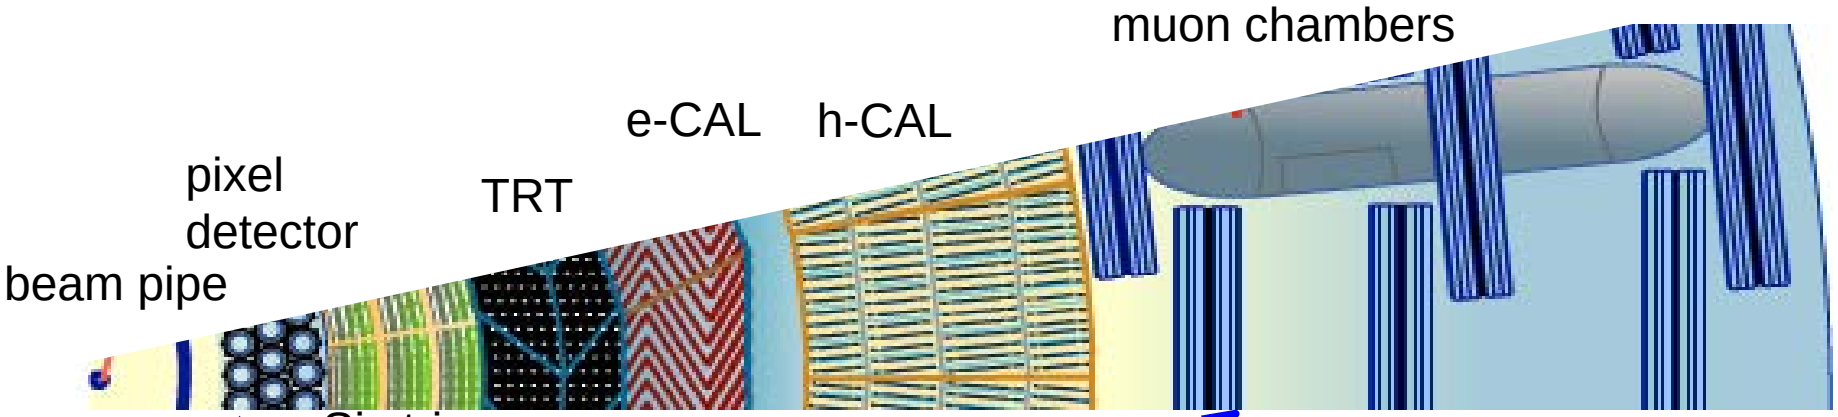

Si strip
tracker

FE-I3 pix det
28 000 chips
80 M segments
1.7 m^2 Si sensor

ABCD Si det
50 000 chips
6 M segments
60 m^2 Si sensor

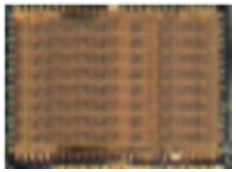

ASDBLR TRT det
38 000 chips

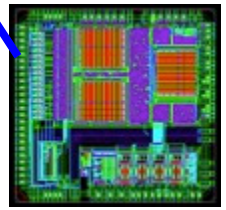

DTMROC TRT det
19 000 chips

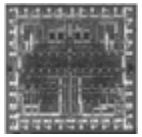

ASD muon det
148 000 chips

Chips to scale 1 cm

Total ATLAS
100 million sensor cells
appr. 800 000 chips
majority ASICs

Some of the ASICs in CMS

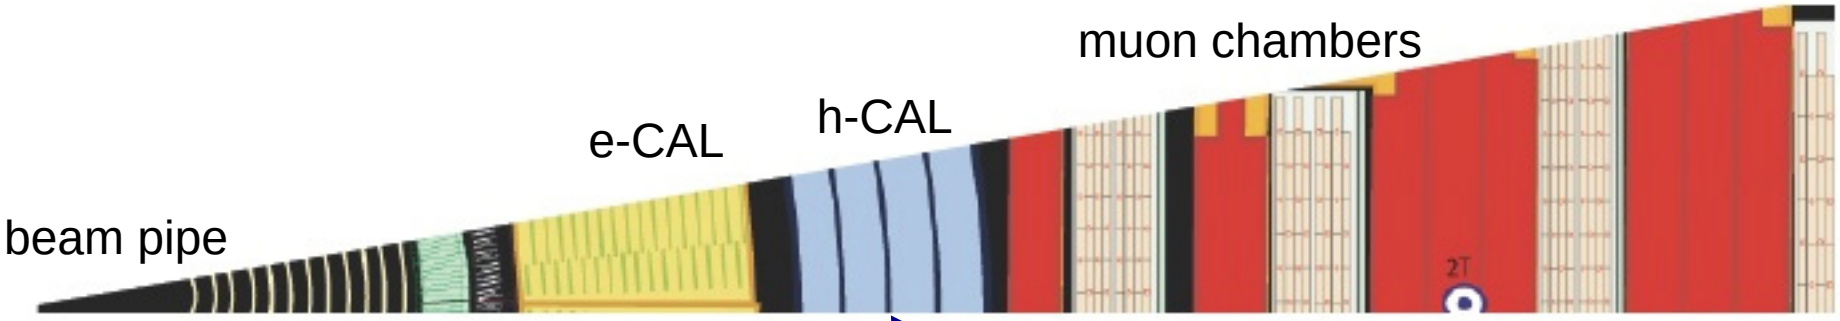

pixel detector

Si strip tracker

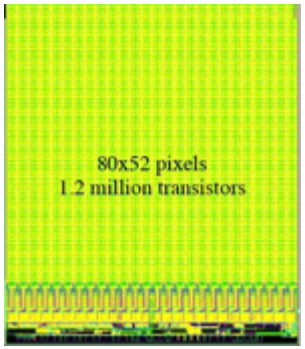

PSI46 pix det
16 800 chips
66 M segments
1 m^2 Si sensor

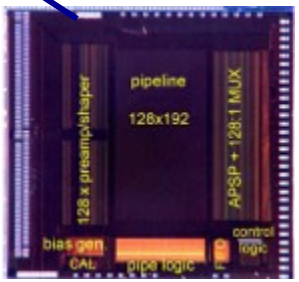

APV25 Si det
110 000 chips
9.3 M segments
198 m^2 Si sensor

QIE8 calorimeter
220 400 chips

MAD muon det
181 000 chips
25 000 m^2 gas-filled

Chips to scale 1 cm

Total CMS
appr. 1 million chips
of which 700 000 ASICs

Animation of bunch crossings and timing

screenshot

~70 sec

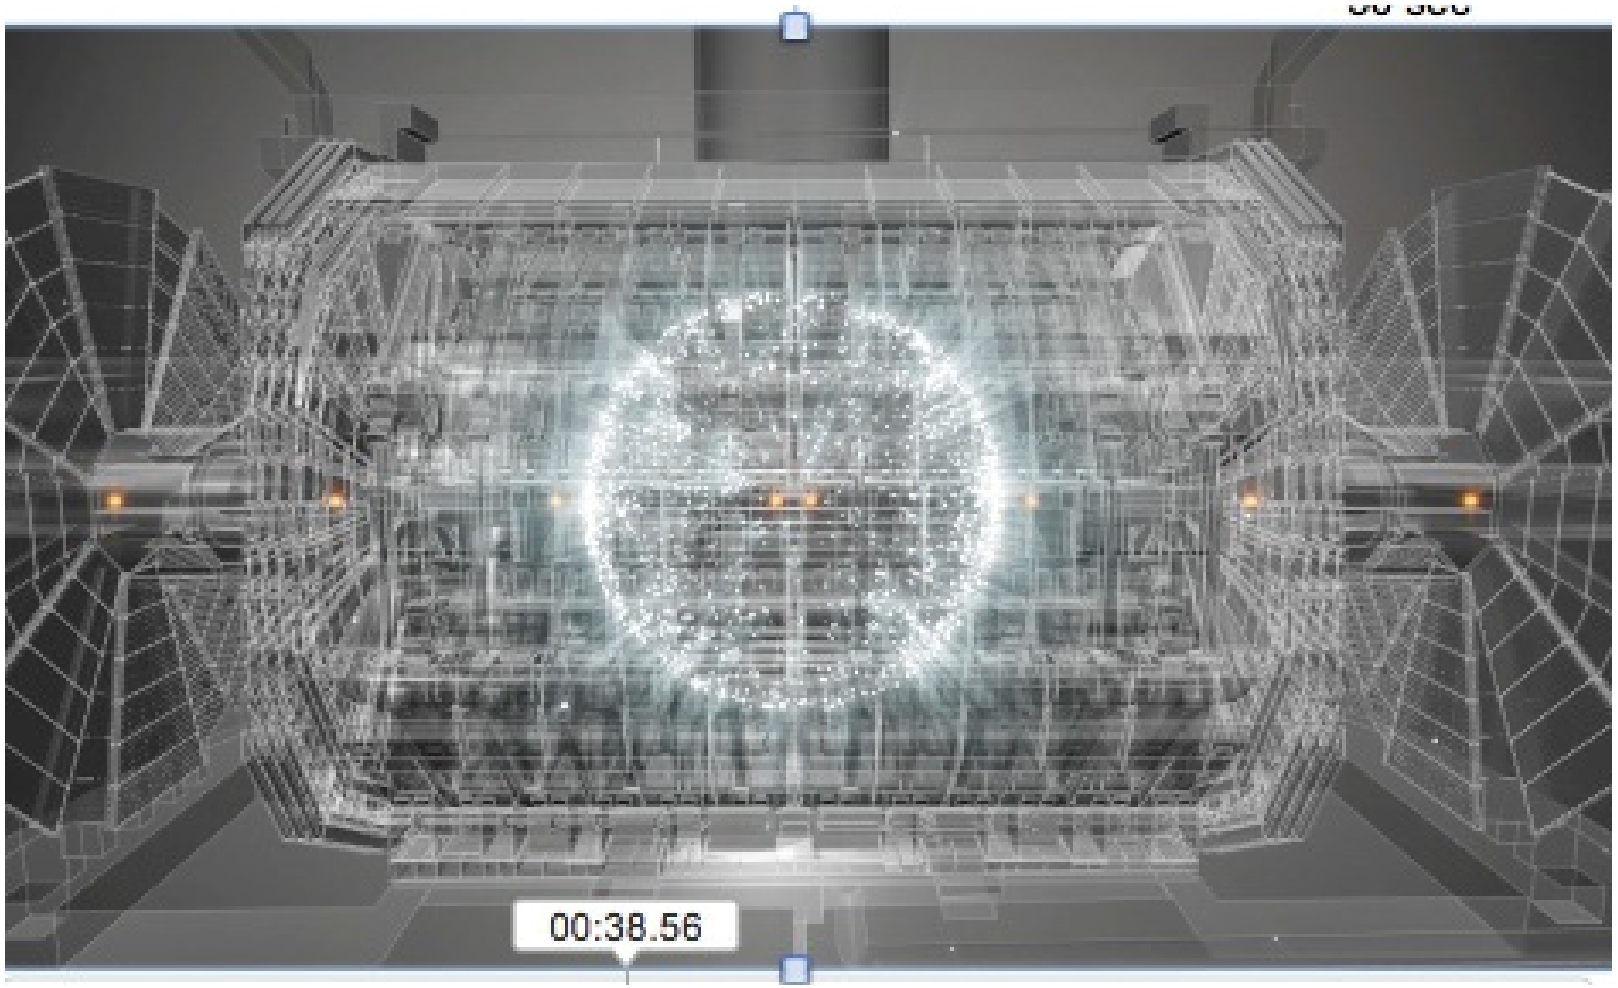

Si Microstrip Sensors and CMOS Readout Chips

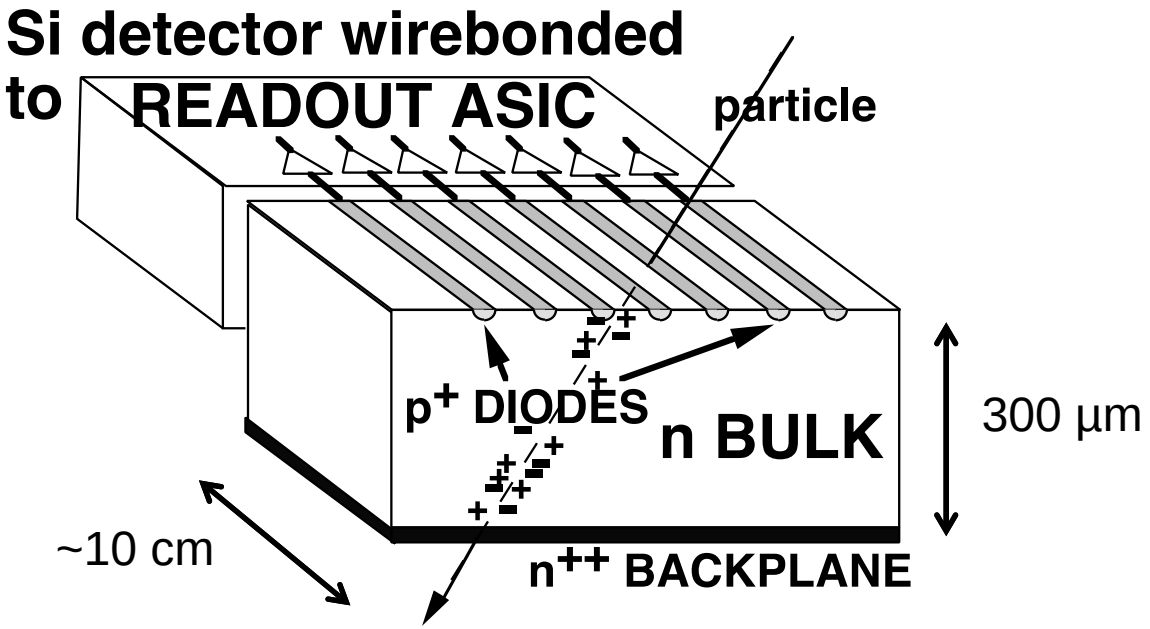

Fully depleted Si Sensor
divided in parallel strips:
“Microstrip Detector”

Individual readout chain
for each segment
Signal ~20 000 e-h pairs

Si Microstrip Sensors and CMOS Readout Chips

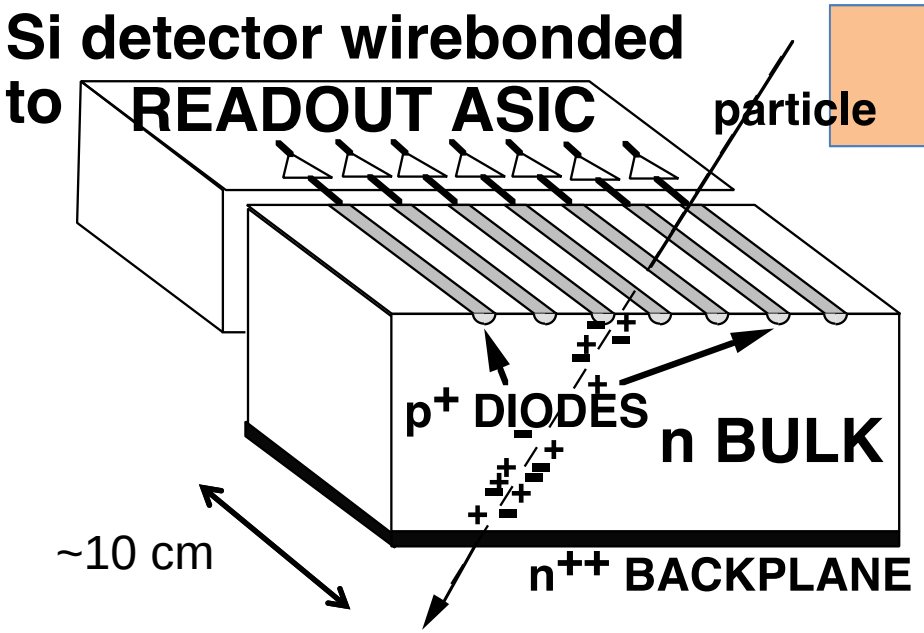

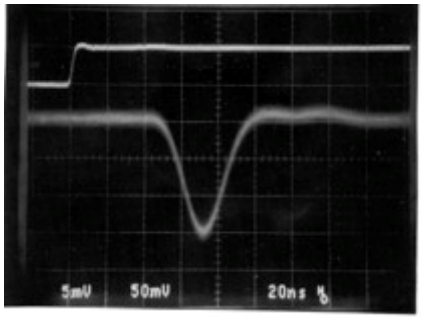

Typical signal
one particle

Fully depleted Si Sensor
divided in parallel strips:
“Microstrip Detector”

Individual readout chain
for each segment
Signal ~20 000 e-h pairs

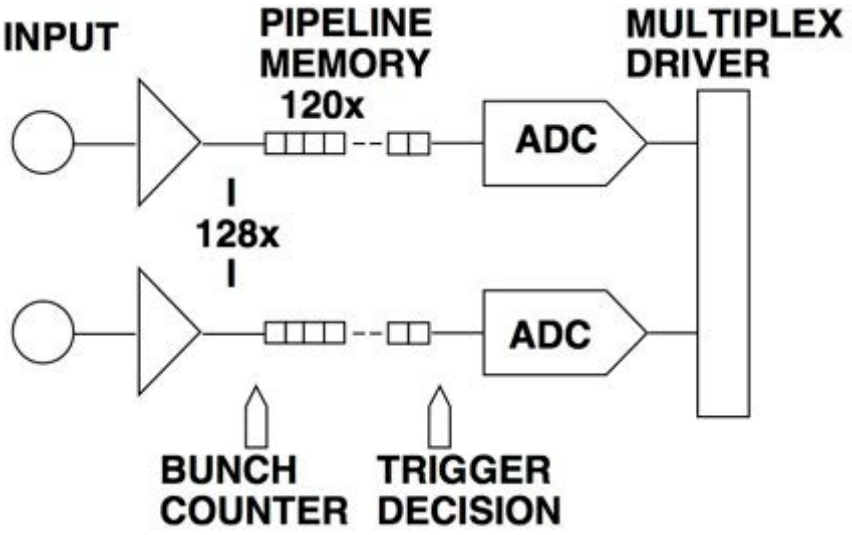

Analog Signal Processing for Si Microstrip Detectors

Example

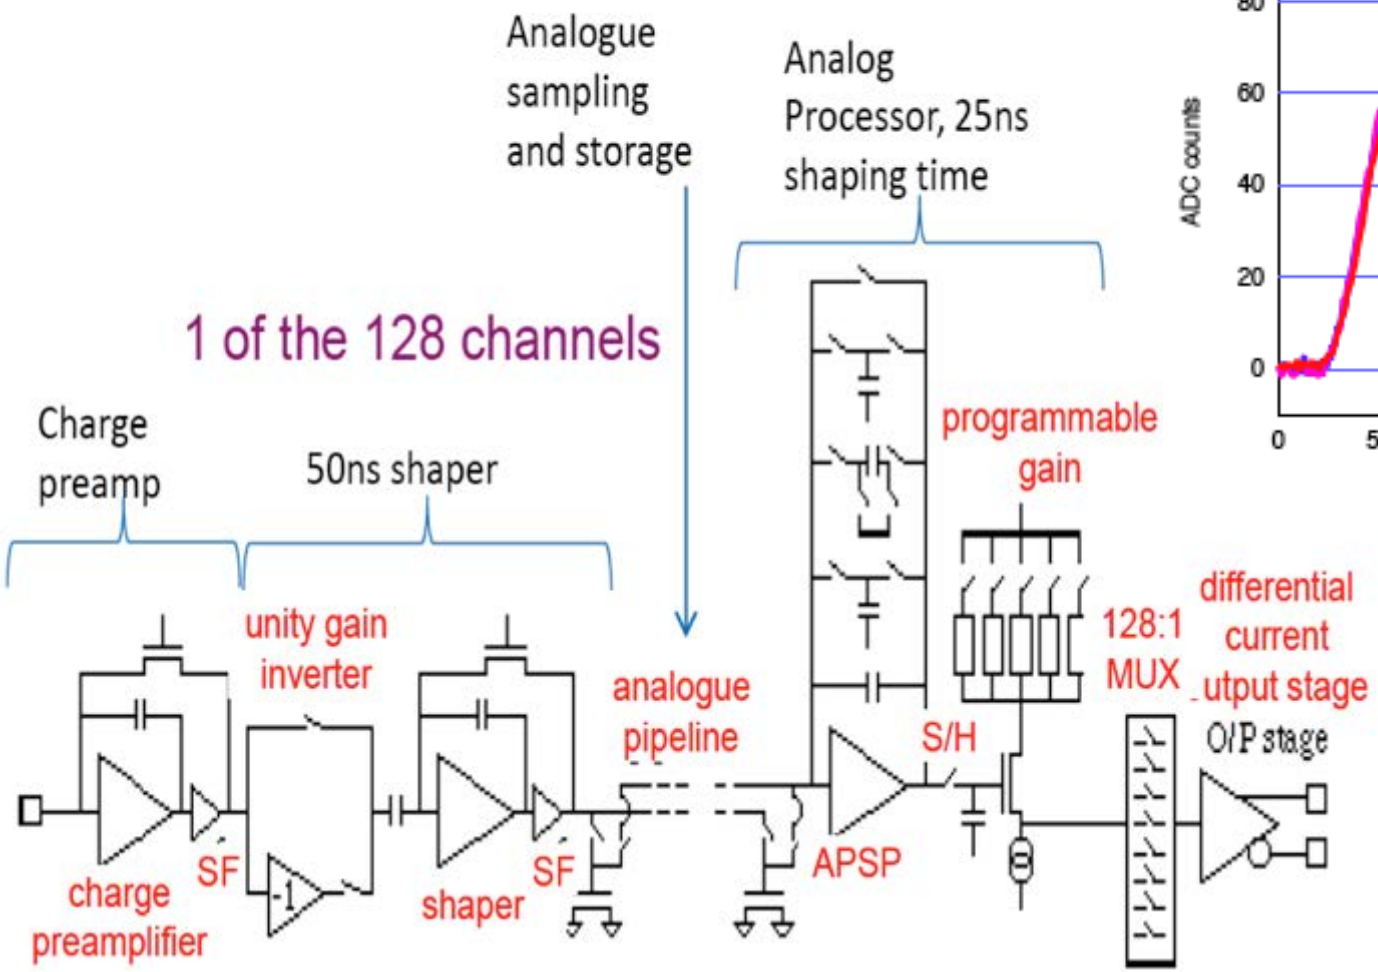

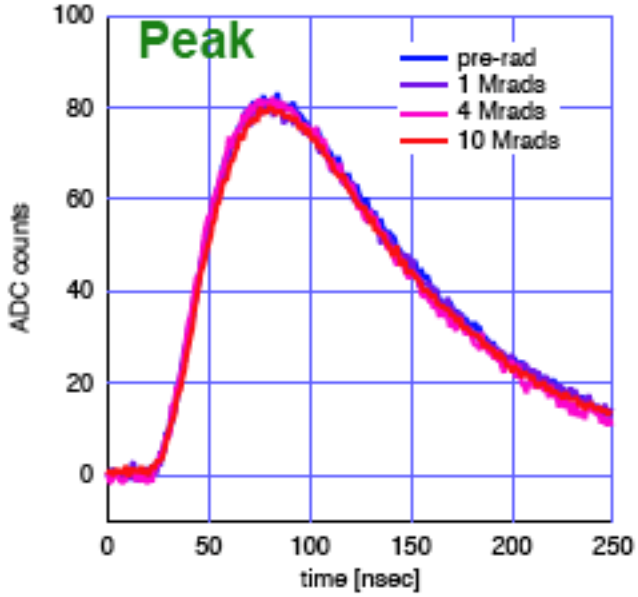

Single channel
CMS APV25 chip
M. Raymond
Imperial College

Characteristics of Readout Chips for Si Microstrip

Amplifier low noise (1500 e^- equivalent)
low power (~mW per sensor element)

Rate capability 40MHz: <20ns time-stamping

Many parallel elements: matching issues

On-chip memory (~4μs) for each element
analog or digital >120 cells

Comparator or ADC on-chip

Radiation hardness: MGy total dose & SEU

ATLAS Si Detector Readout “ABCN25”

Example

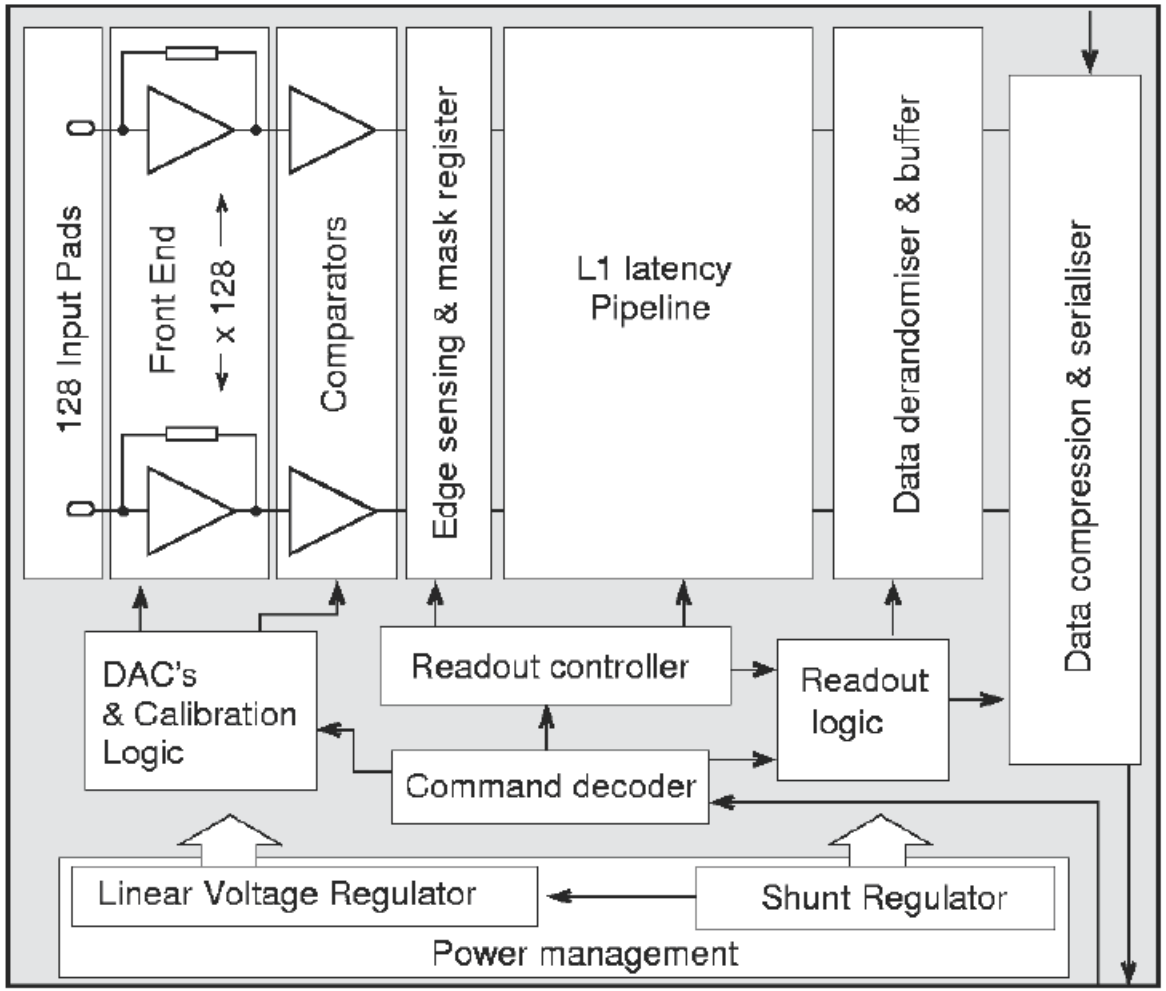

128 parallel channels
Amplifier + Comparator

Binary Pipeline Memory

ATLAS Si Detector Readout “ABCN25”

Example

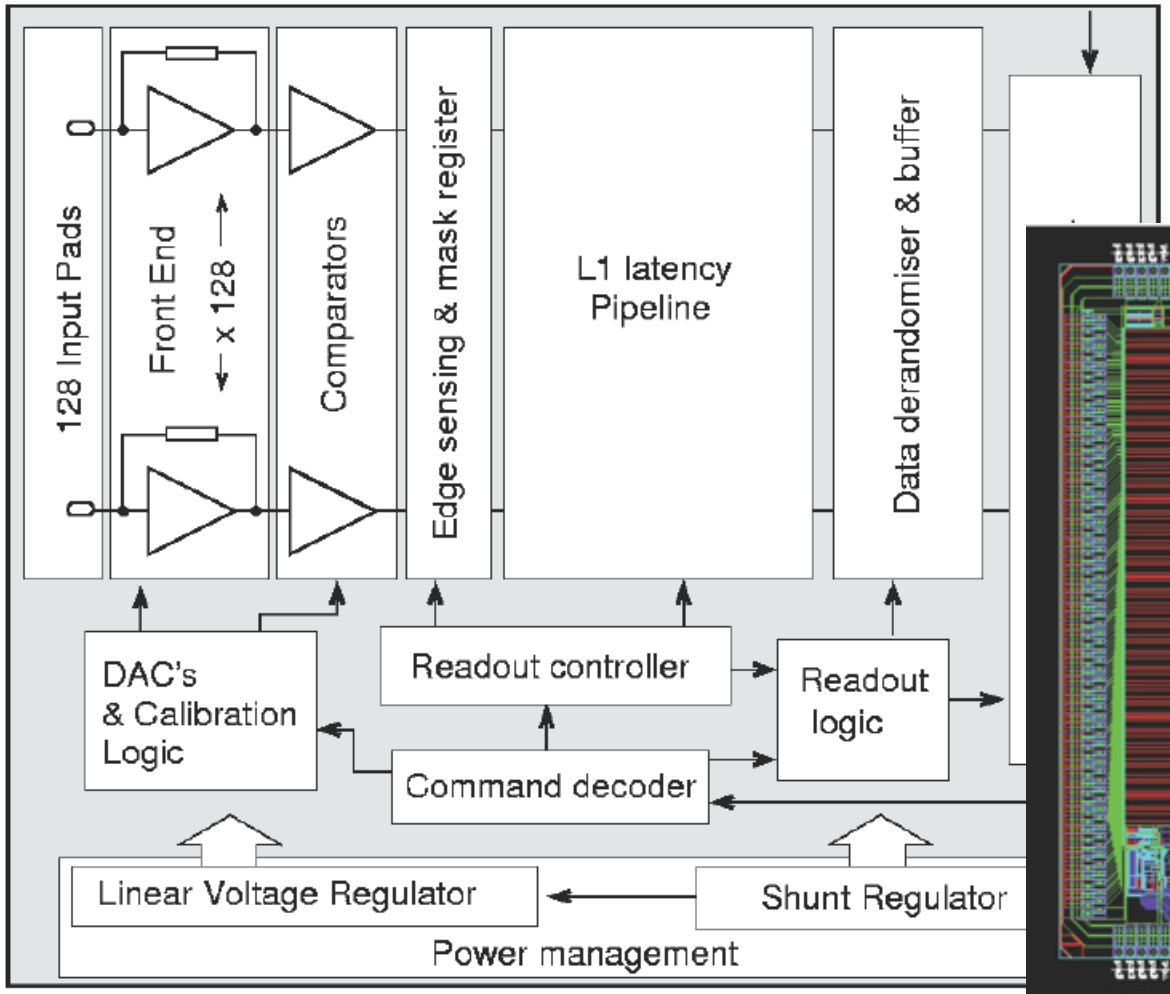

128 parallel channels
Amplifier + Comparator

Binary Pipeline Memory

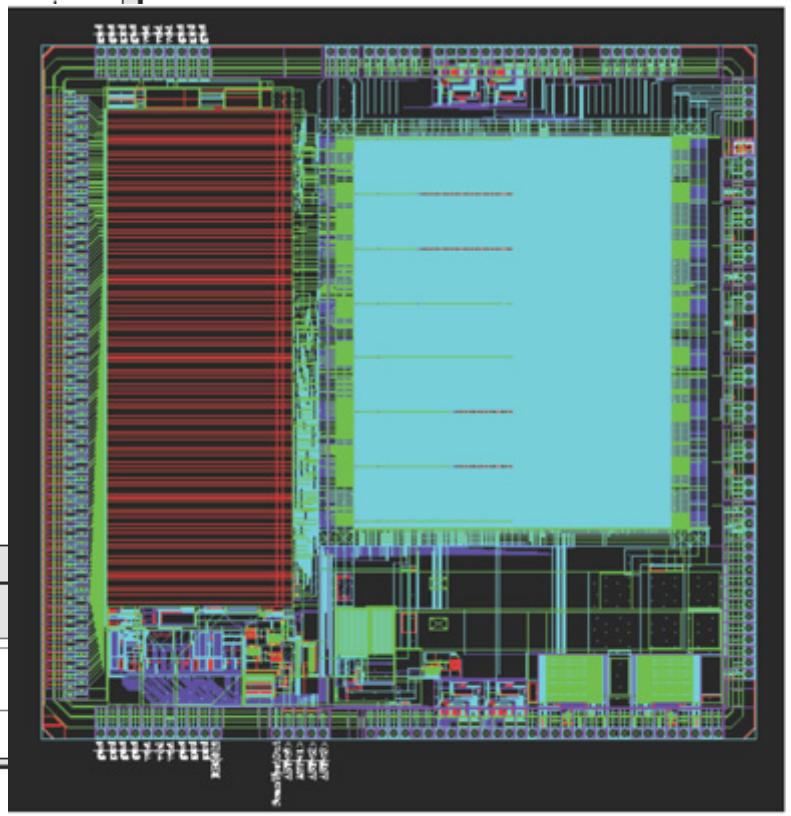

7.7 x 7.5 mm²

Example of Evolution: Data Storage

~1985 parallel-serial CCD analog pipeline

1992 analog storage on feedback capacitors

2002 comparator first, binary pipeline storage

iteration “ABCD” actually installed in ATLAS

2004 SRAM allows variable retention times

 construction & installation

Silicon Microstrips in Particle Physics Experiments

MOS technology for position measurements of ionizing particles

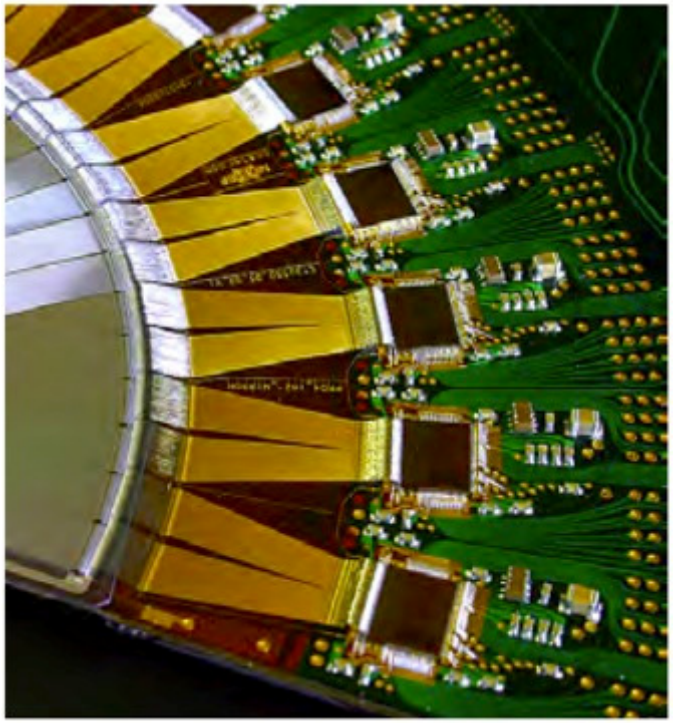

Si sensor connected to ASICs

in “small” LHCb experiment

Barrel of Si sensor modules
ATLAS employs 60m^2 in 3 layers

source: CERN

Radiation Hardness: e.g. SRAM on-chip

an example

Radiation hardness “by layout and by design”

- using standard CMOS, not by special technology
- we developed own expertise
- extensive contacts with world specialists

Larger than minimum size, larger than minimum current

Transistor layout avoids leakage current under field oxide

Triple redundancy on memory pointers and timing

- Single Event Upset errors in data are tolerated

Enclosed n-Transistor is Radiation Tolerant

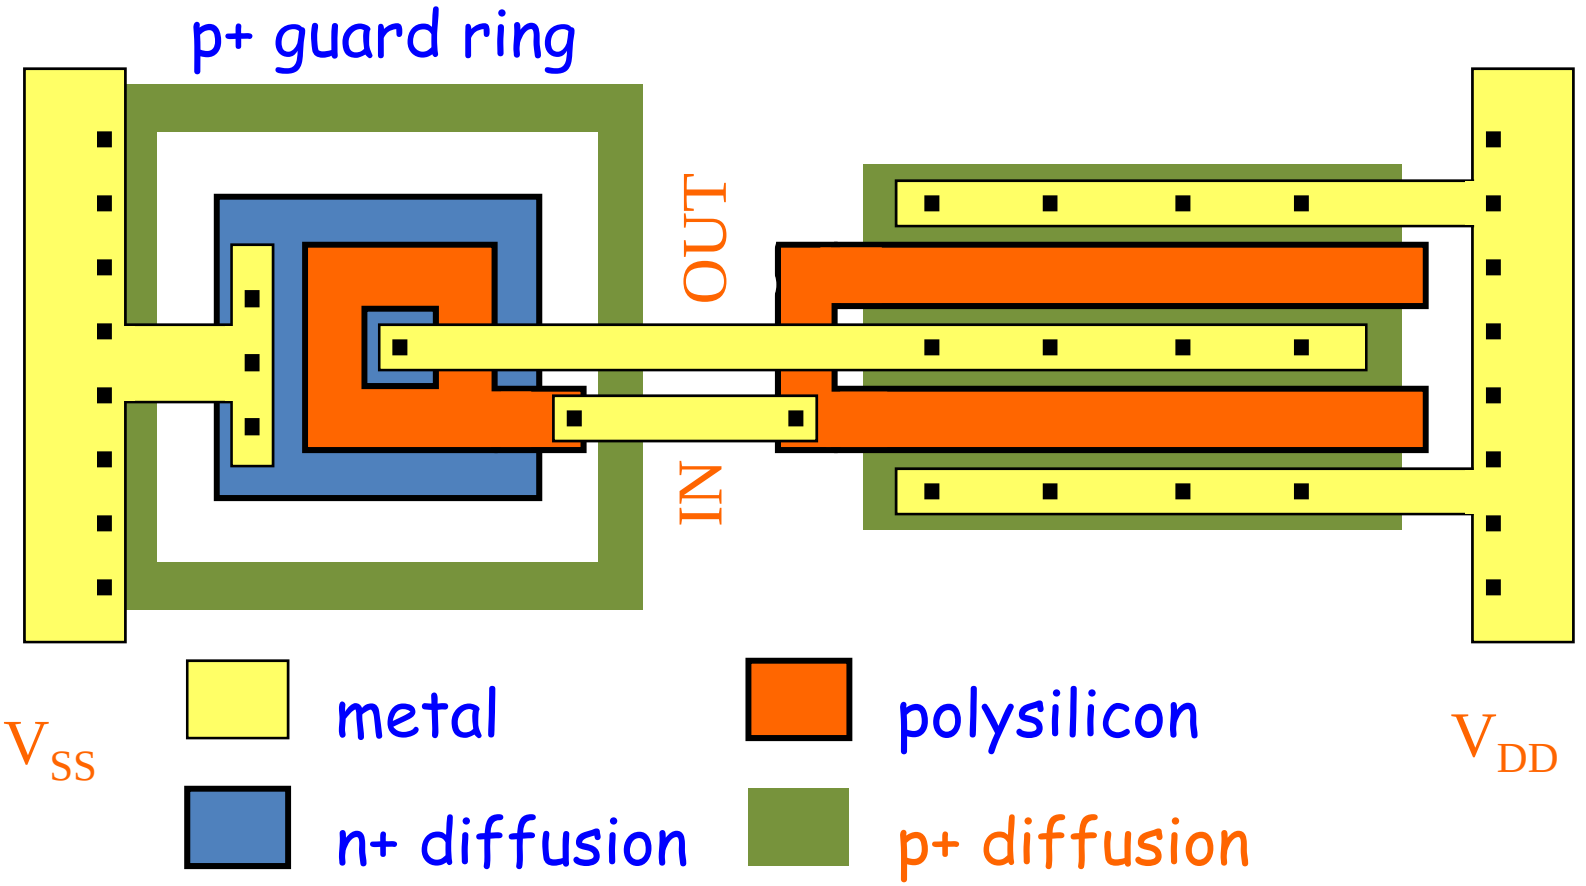

need to study physics of radiation effects in devices

Enclosed n-Transistor is Radiation Tolerant

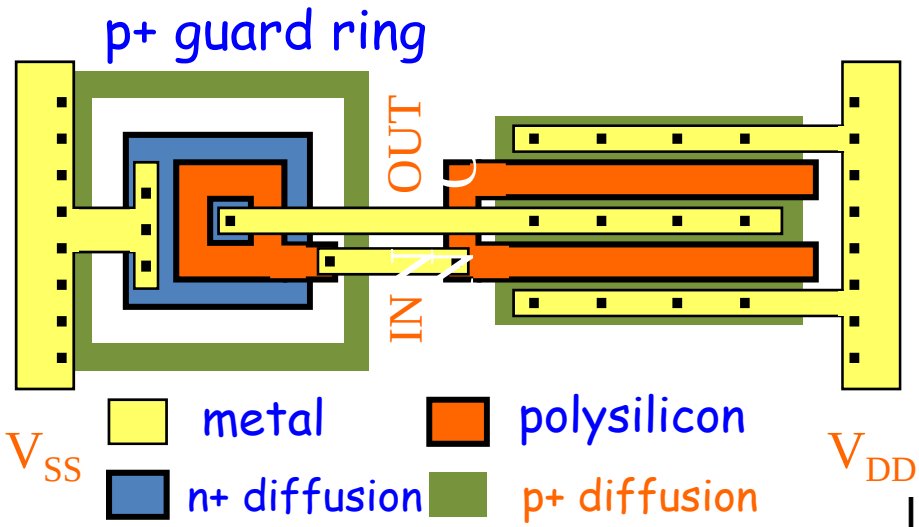

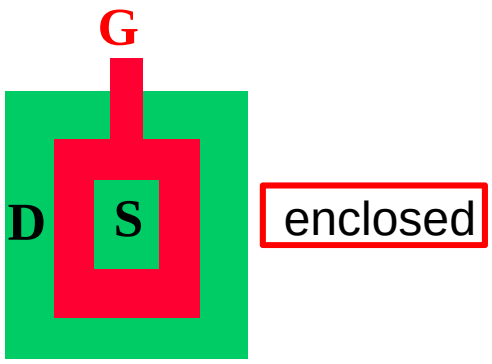

Irradiation 3Mrad, 136 Mrad (1.36 MGy)

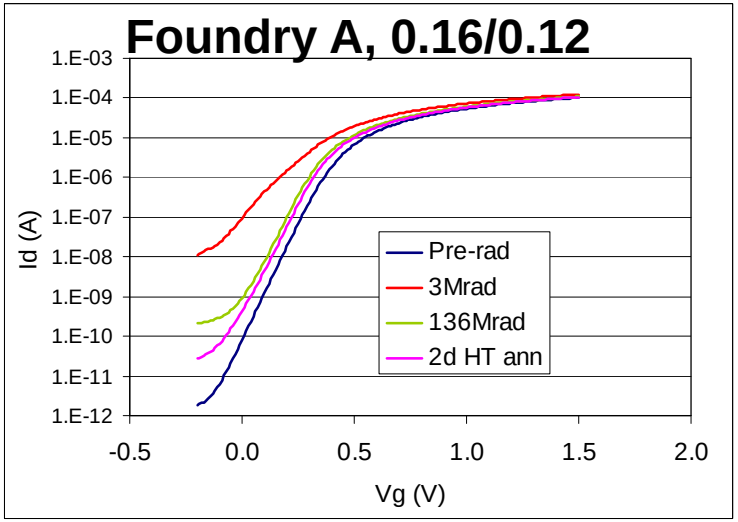

Id- Vg min size, linear layout

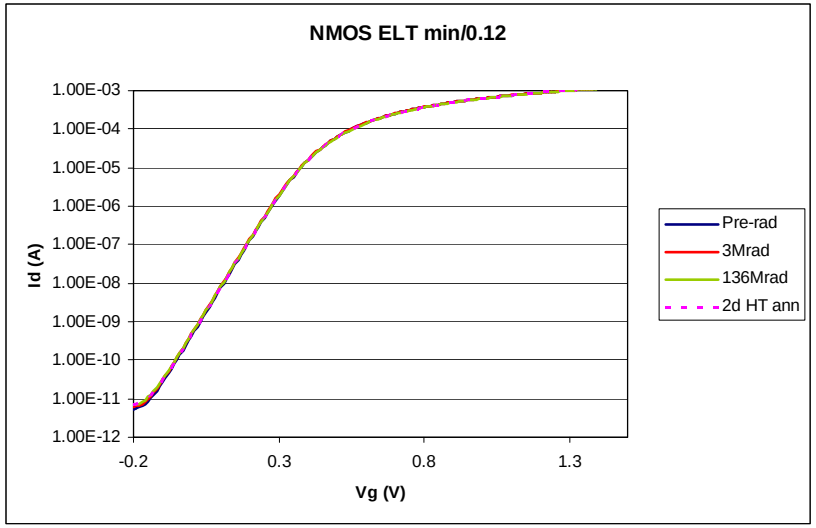

Id- Vg min size, enclosed layout

Silicon Pixels in Particle Physics Experiments

Developed for high densities and high rates of ionizing particles in LHC

Segmented Si sensors achieve best possible particle tracking precison
(only photographic emulsion can do better, but is not electronic)

ATLAS pixel FE-I3 basic cell

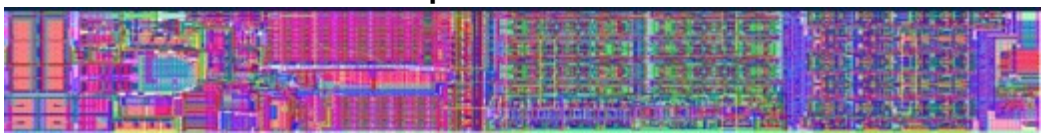

50um x 400µm
0.25µm CMOS

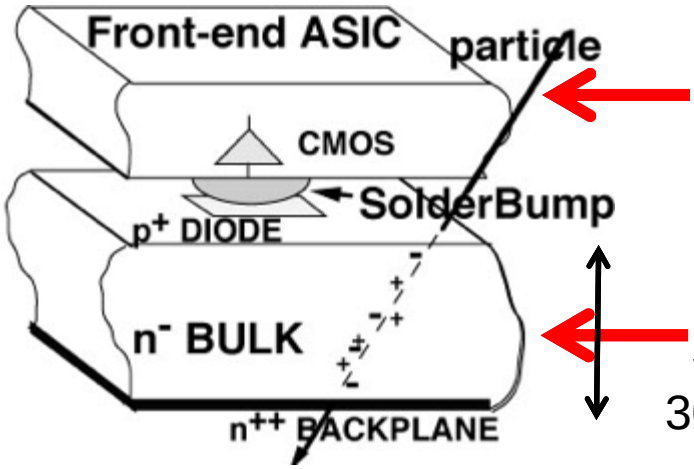

Readout matrix

Sensor matrix
300 µm

Si sensor matrix is connected
by matrix of solder bonds to ASIC

Spin-Off Project/Demo
Timepix basic cell

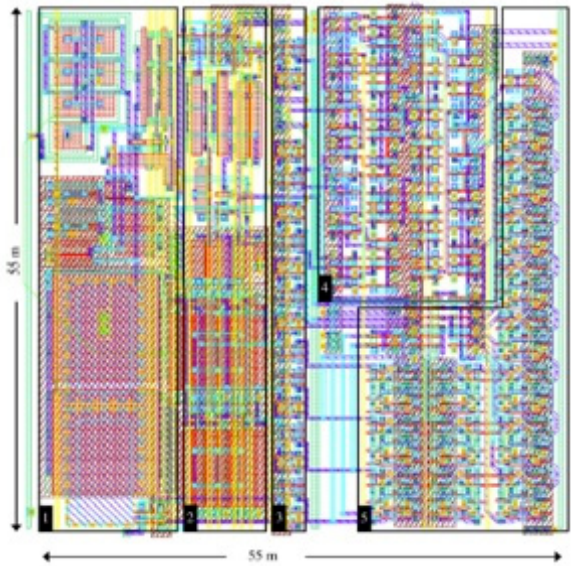

55 µm

55 µm

Layout of a pixel cell
55umx55um, 500 transistors
14mm chip has matrix of 256x256

Block Diagram of Circuit in each Pixel ATLAS

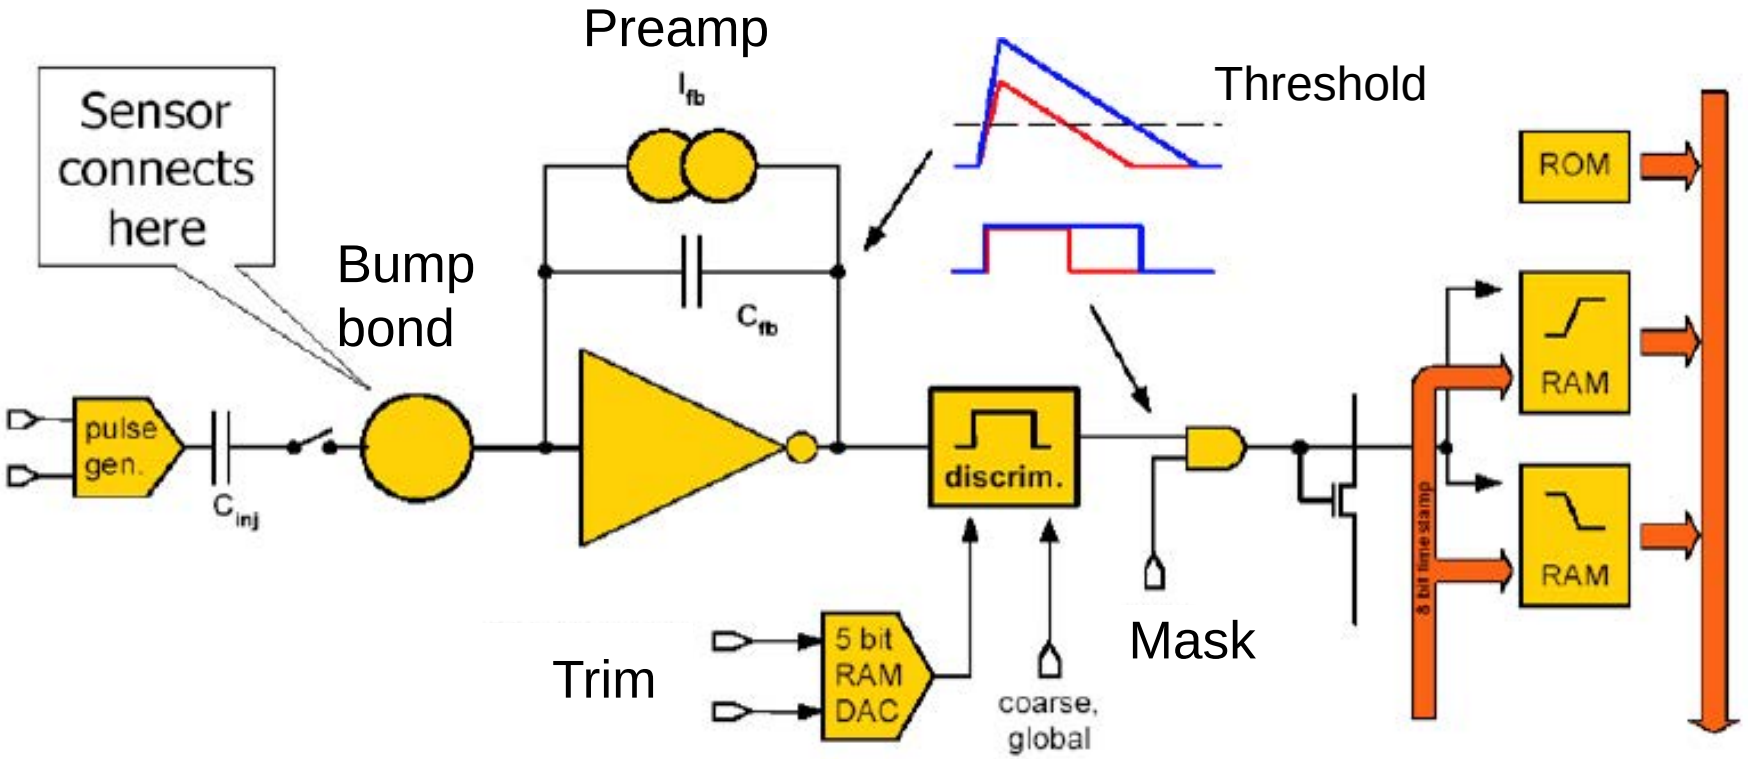

after M. Garcia, LBNL Berkeley

Pixel-readout Chip “FE-I3” in ATLAS

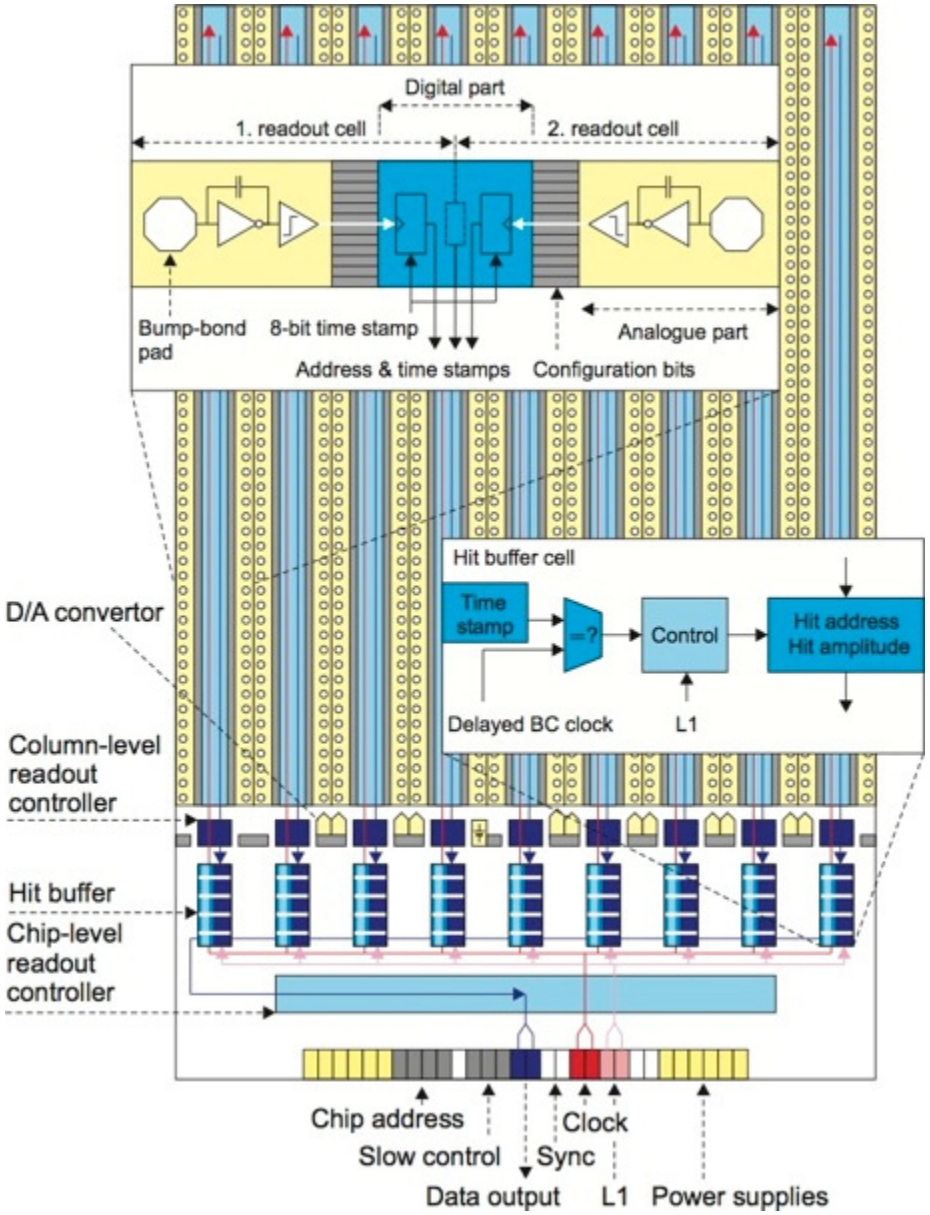

11 mm

7.4 mm

18 x 160 pixels of 50µmx400µm

ATLAS Inner Si Pixel Layer

source: CERN- ATLAS

Collision Event at 7 TeV with 2 Pile Up Vertices

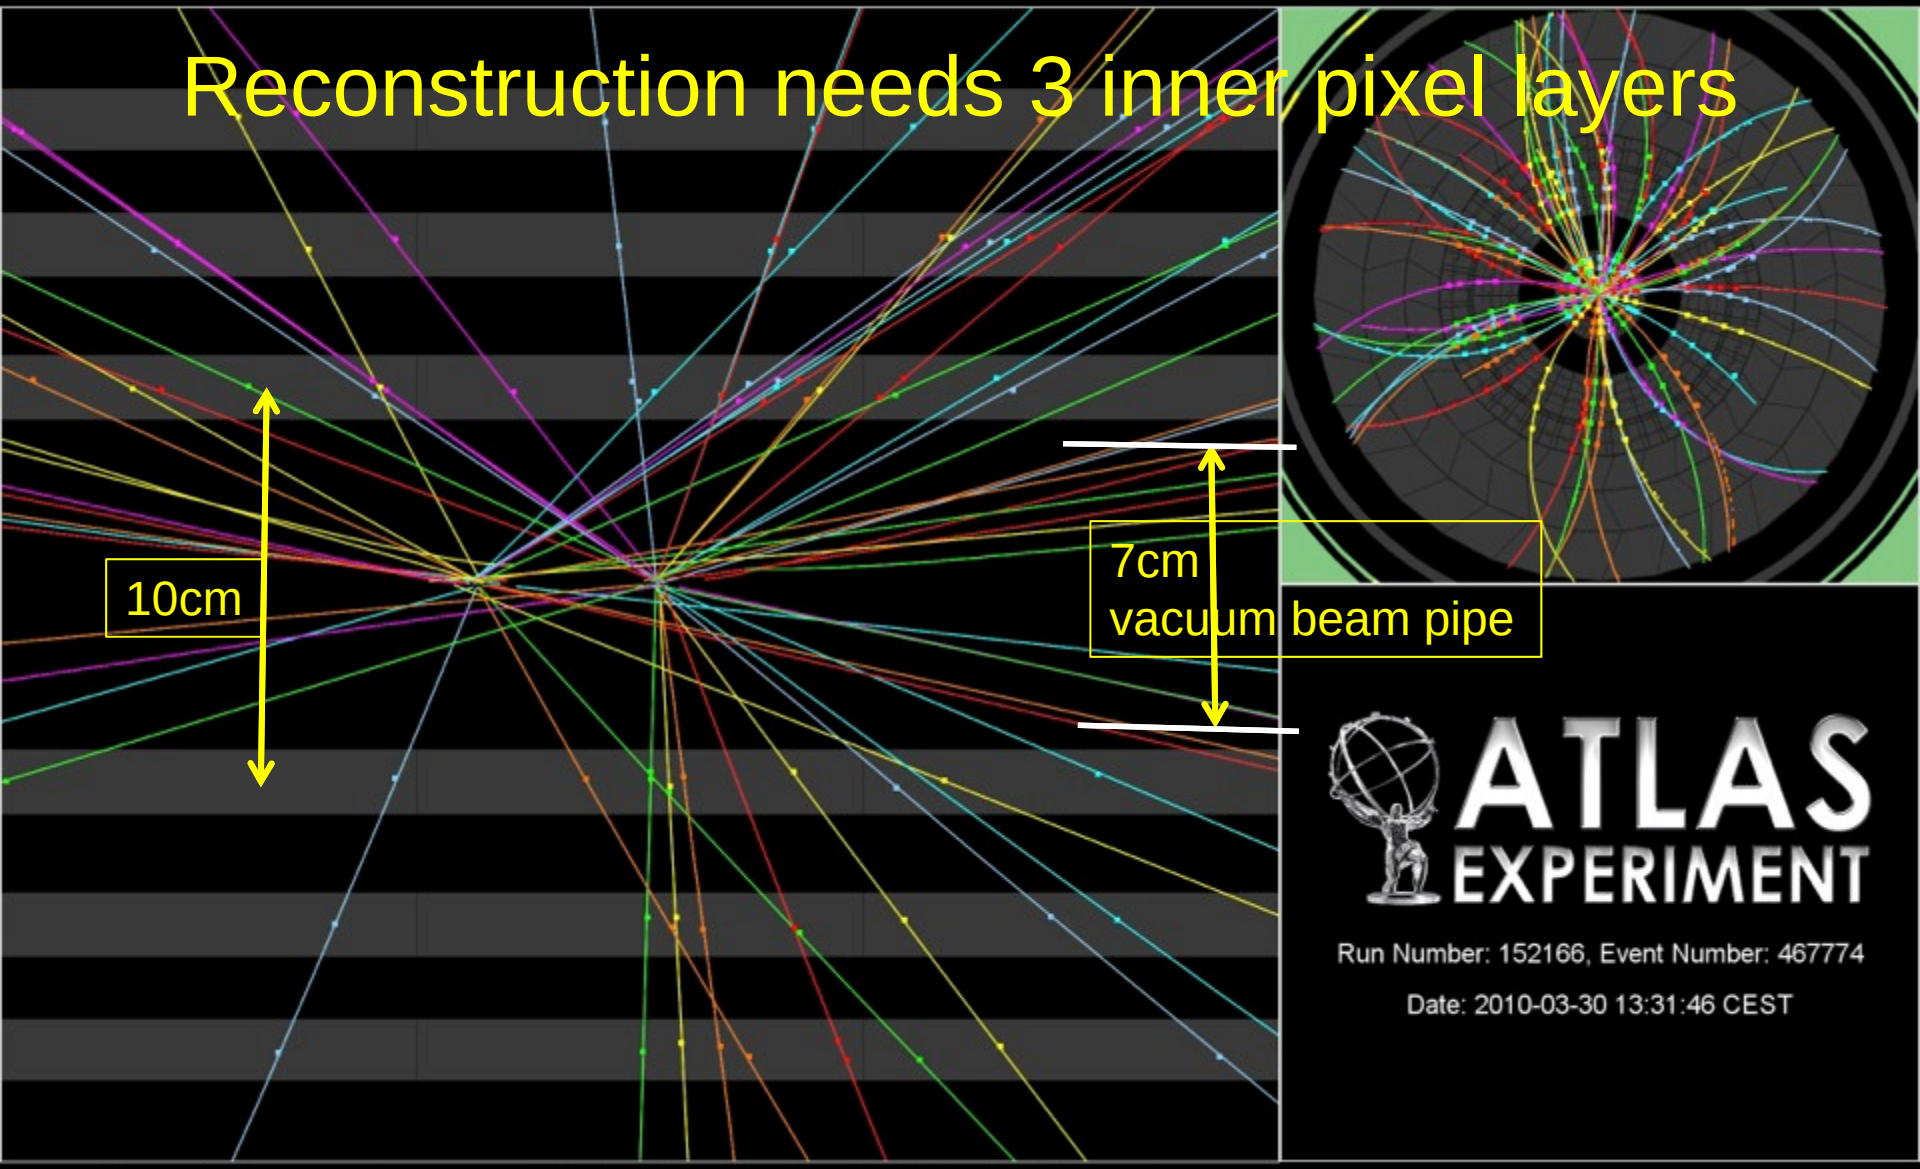

http://atlas.web.cern.ch/Atlas/public/EVTDISPLAY/events.html

source: CERN- ATLAS

Intermezzo: Radiation Quanta in Pixels Portable Detector as USB Stick

See also Monday Afternoon Demo Session paper 1.3

Single Quanta around Us

60s Exposure at ground level

60s Exposure in plane at 24 000 feet

6s Exposure with old wristwatch (radium)

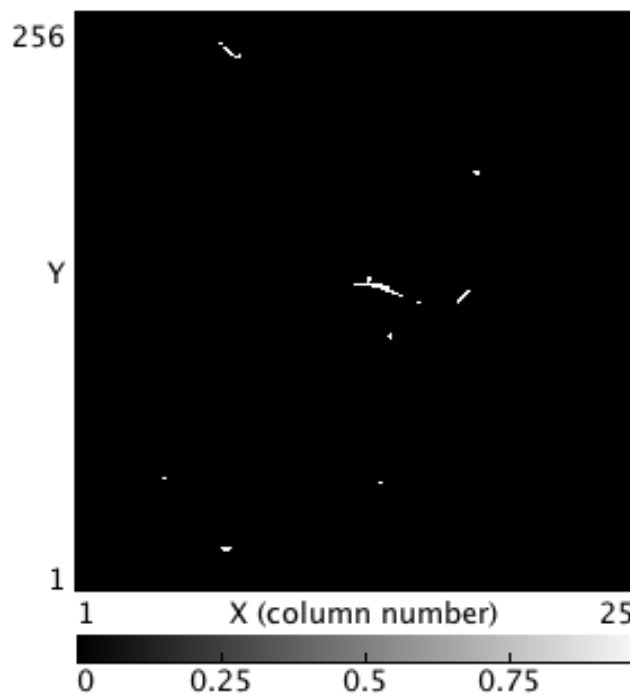

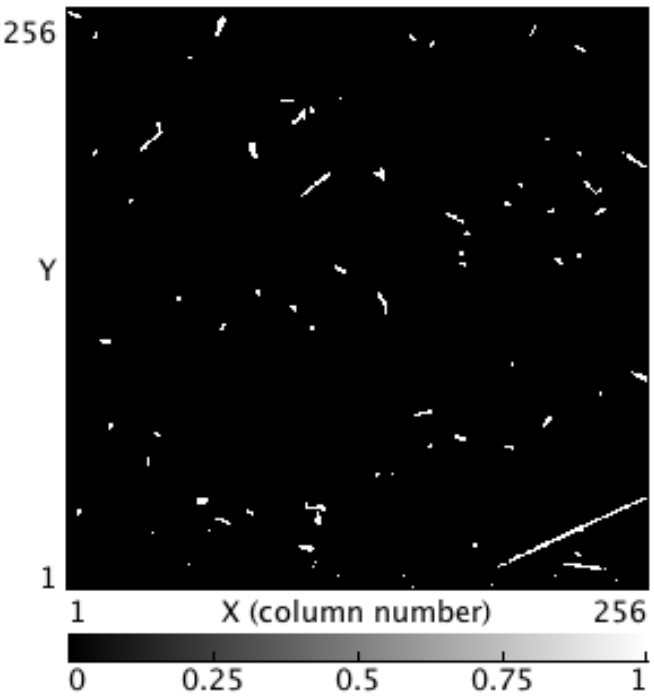

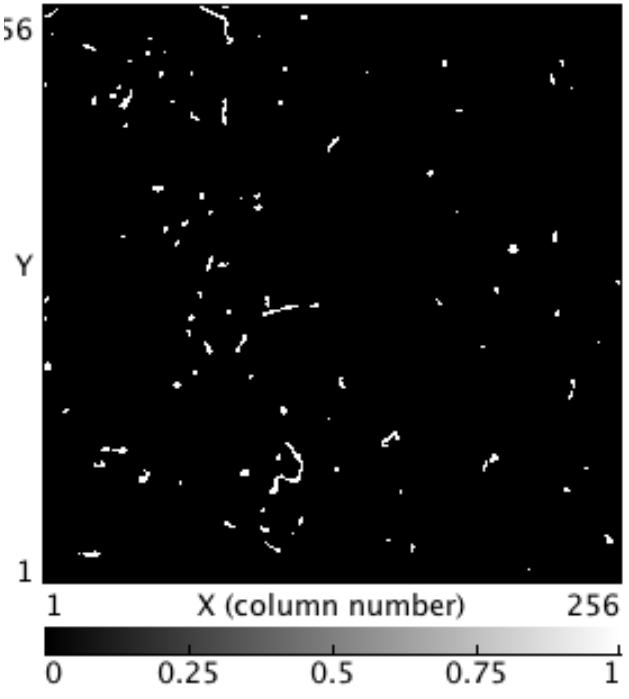

Dose levels well within safety limits

Data from Continuous Frame Rate at 40MHz

in the Large Hadron Collider LHC
4 experiments generate each
~1Mb per 25ns

Intelligent Filtering, yet Big Data
Dedicated Cloud System for the Analysis:
Worldwide LHC Computing Grid
WLCG

Worldwide LHC Computing Grid WLCG

Dedicated Physics “Cloud”

4 Tiers:

0- CERN

1- Main Centers 13

2- Universities 156

3- Users ~8000

Capacity Dec 2013

CPU physical 142

CPU logical 466

Online 0.27 ExaByte

Nearline 0.14 Exabyte

New Data Input 25PB/year

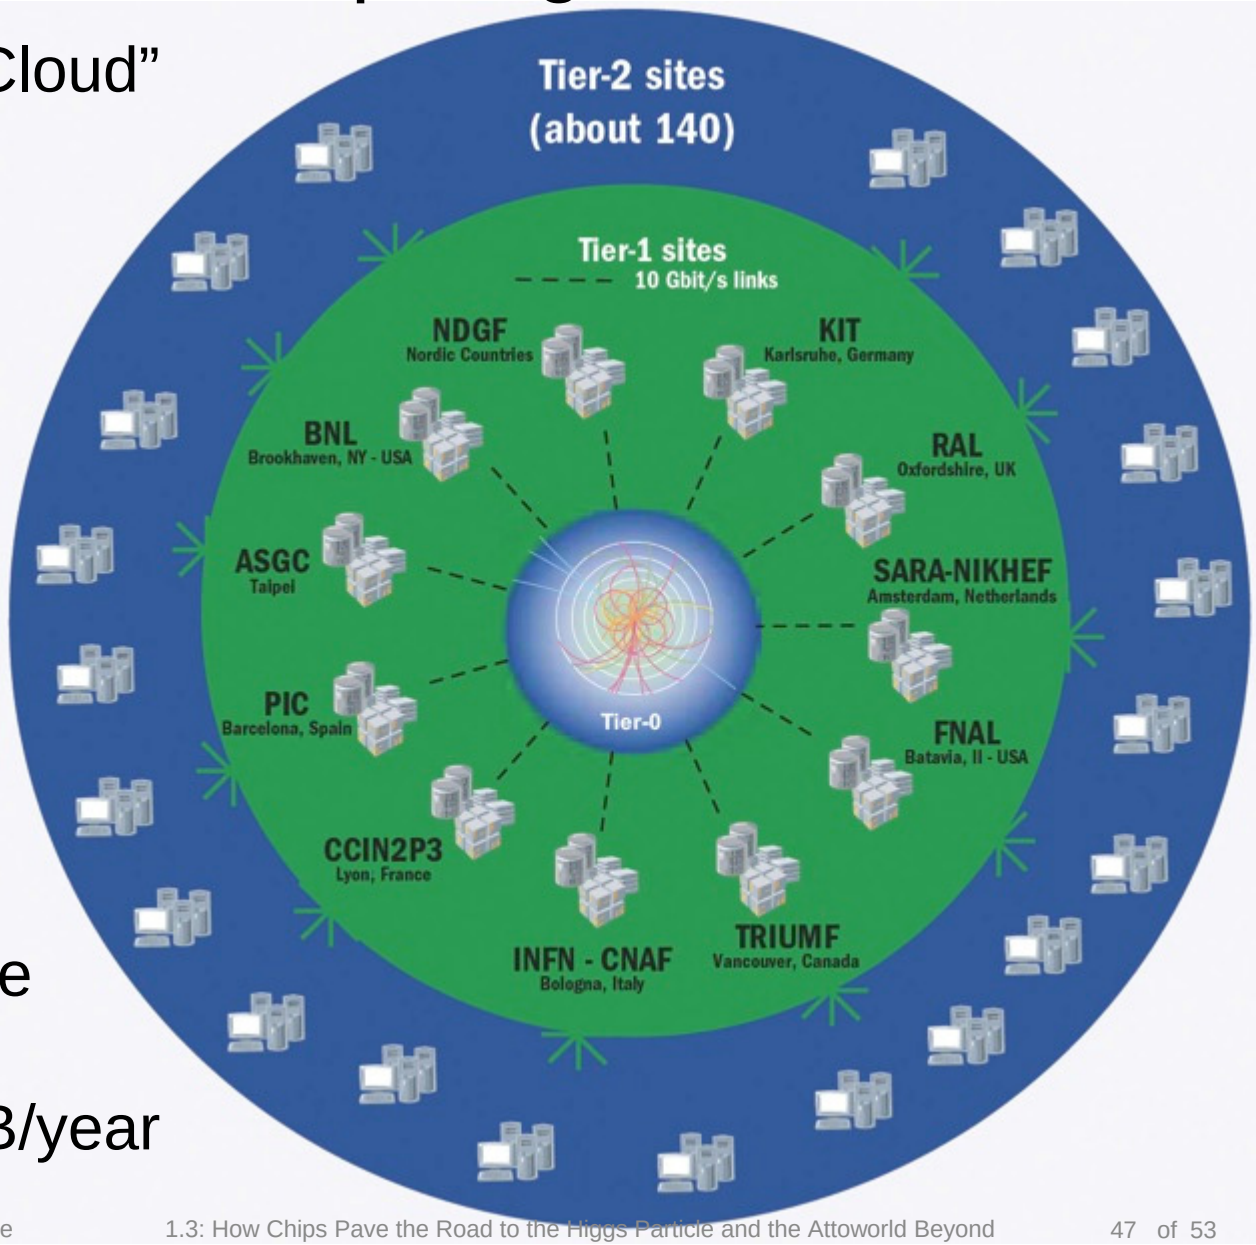

World Wide Web Started at CERN in 1989

HyperText Transmission Protocol http
by Tim Berners-Lee with support of Robert Cailliau

First server
“NEXT” workstation

source: CERN

ATLAS evolution of real data: Higgs → 2 Photons

screenshot just before Higgs announcement ~25 sec

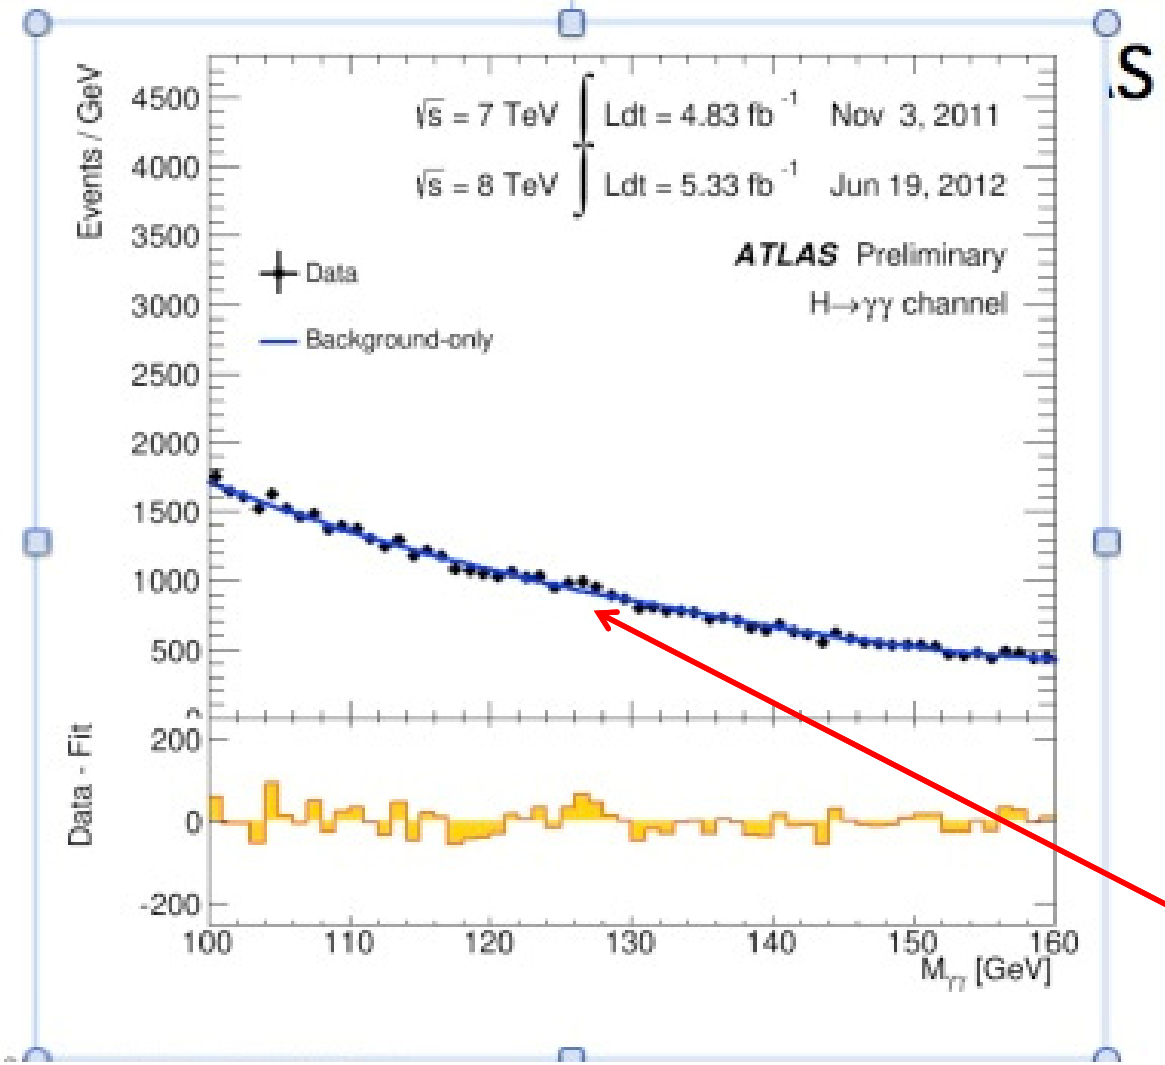

All “events” with 2 energetic photons emerging from the primary vertex

Ordered by sum of measured energy in GeV of all decay products in each interaction

An excess of events is observed around 126 GeV

just OK at 4 July 2012

ATLAS

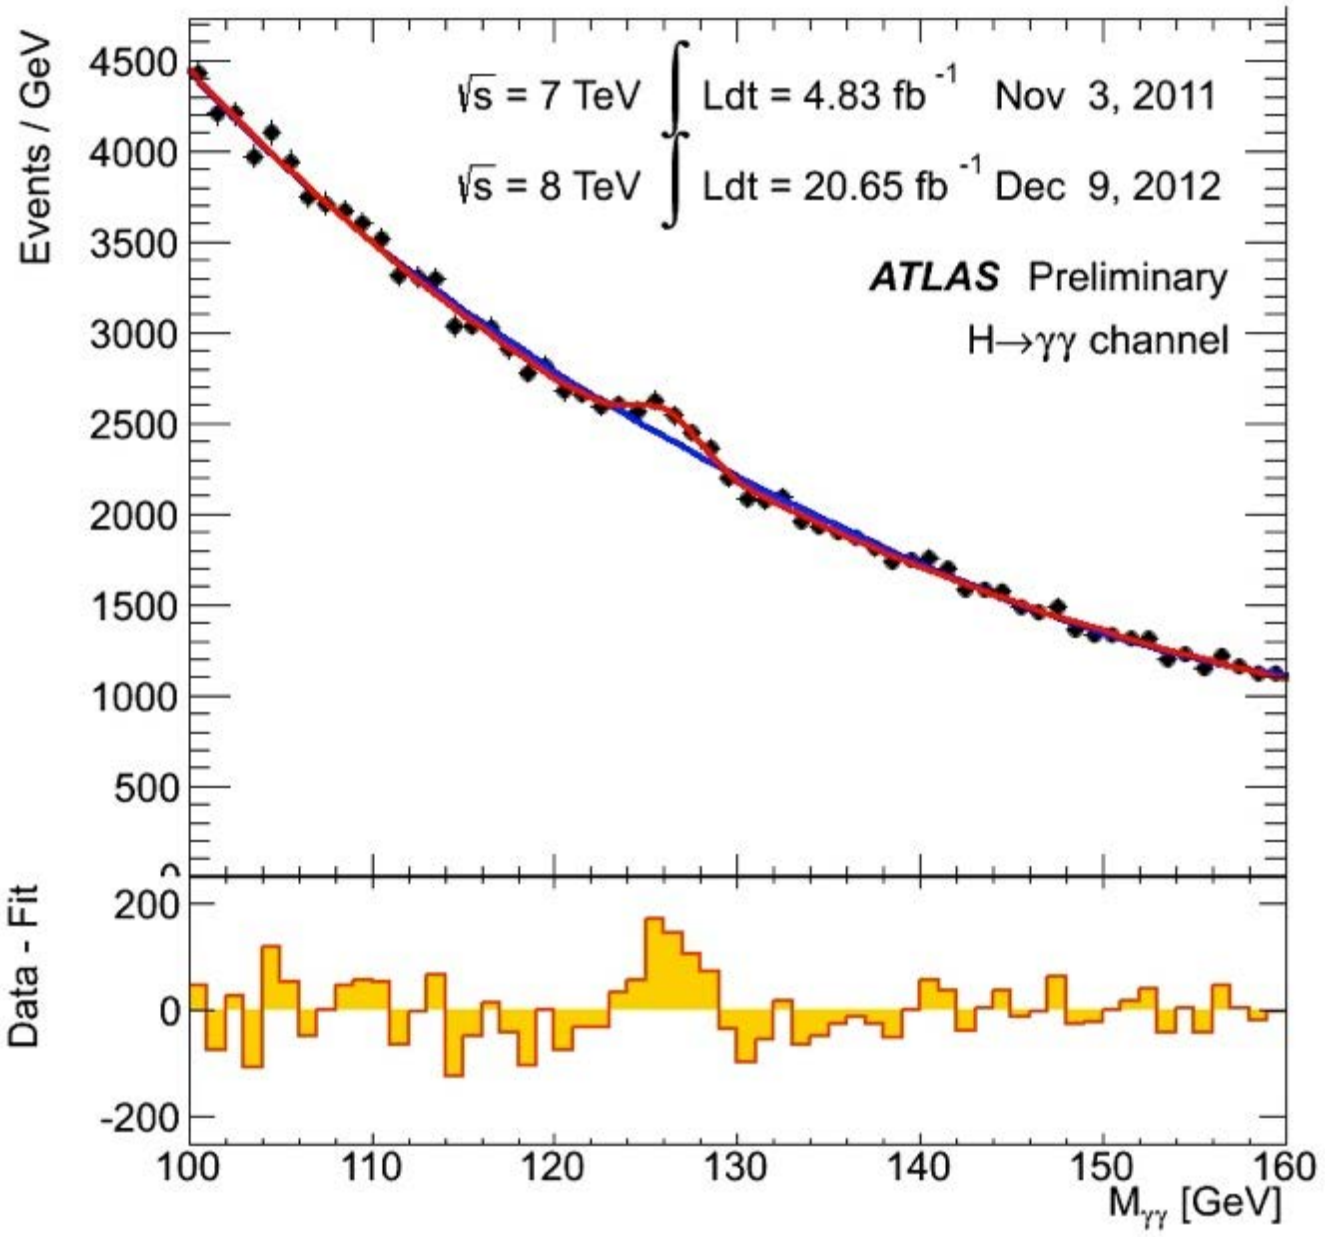

All “events” with
2 energetic photons
emerging from the
primary vertex

Ordered by sum of
measured energy
in GeV of all
decay products
in each interaction

An excess of events
is observed around
126 GeV

just OK at
4 July 2012

source: CERN- ATLAS

Higgs Boson: What is it All About?

4 July 2012 Discovery of a "Higgs-like" particle

Summer 2013 Confirmation of a Higgs-boson

December 2013 Higgs-boson decays into Fermions

The implication of this discovery is:
a new, until now unproven, "Higgs" field permeates space, interacts with particles and determines their mass

Chips for Science
Elementary Constituents, Forces and Fields

Emphasis in this talk was on the impact of
silicon chips and sensors in the experiments

Higgs Field now proven
with Higgs-boson as carrier
Theory of "Attoworld" complete (?)

Will any of you imagine future applications in technology?

Outlook and Conclusion

Improvement of LHC experiments

- operation planned until 2035
- more sensors
- more intelligent filtering of data

Discovering new physics

- supersymmetric particles
- dark matter ....

Physics useful for new technology?

- practical: superconductivity
- work with single, energetic quanta
- quantum circuits coming, but slowly

The Next Generation of Networked Experiences

Susie Wee

VP and CTO of Networked Experiences

Cisco Systems

Applications drive the evolution of networks

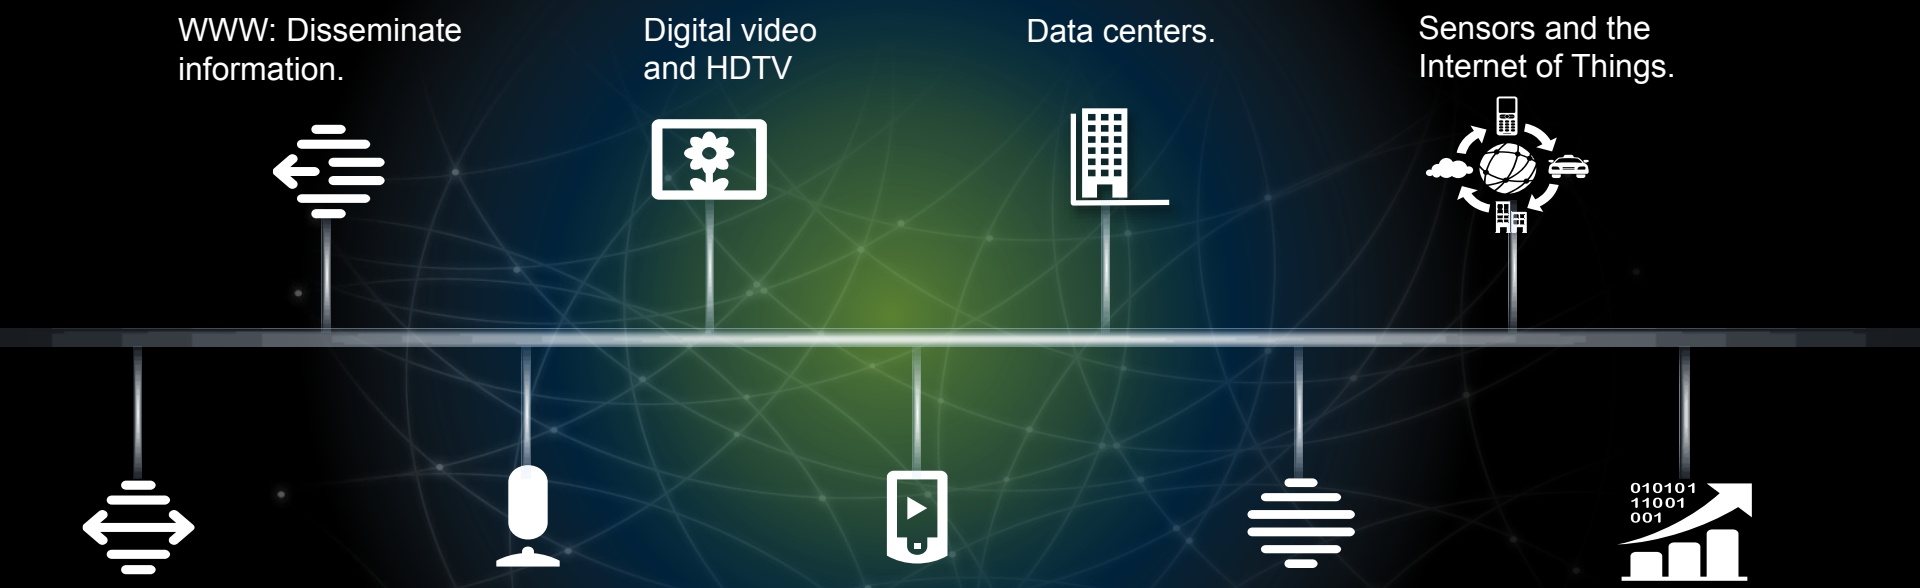

Networked Experiences

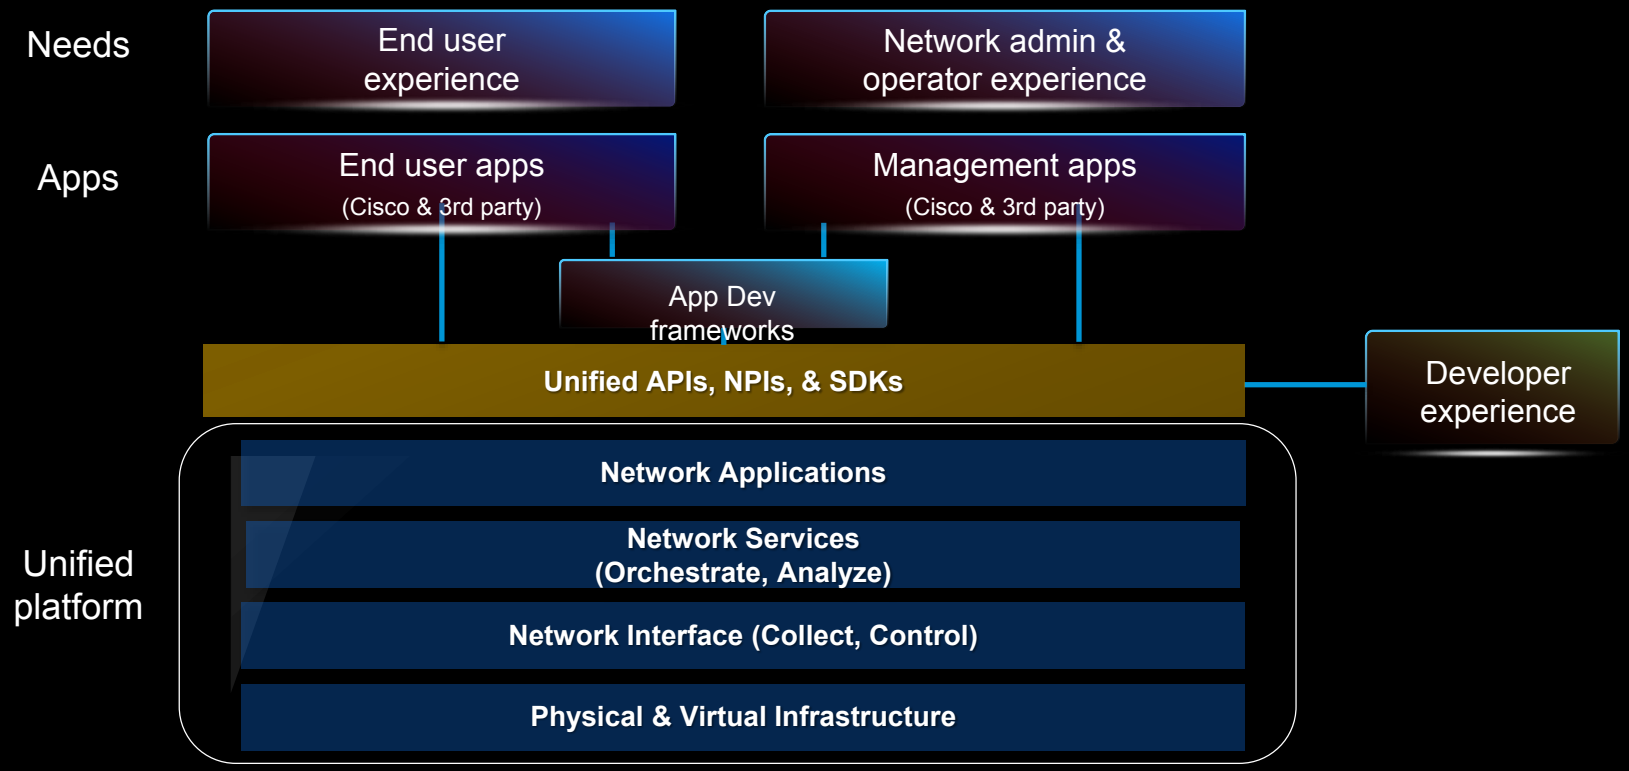

Today's Major Topics

Augmented Collaboration

A next-generation end user experience

Network Function Virtualization (NFV)

The evolution to software-based networking

Software-Defined Networking (SDN)

The next-generation network operator experience

Augmented Collaboration

A Next-Gen End User Experience

Cisco Collaboration Portfolio

When you're mobile

At your desk

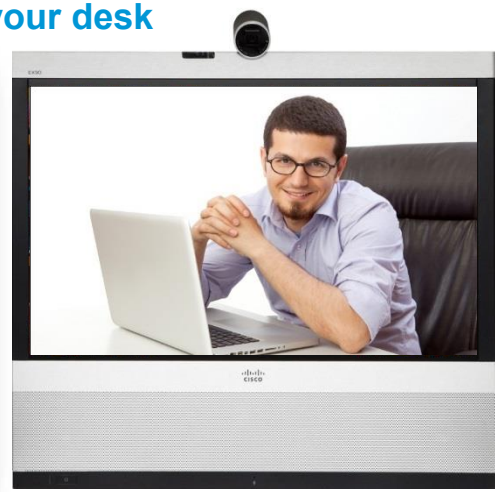

Or in the conference room

Demo Video

Augmented Collaboration: Touch Gestures

1 finger tap

Select

1 finger press & hold

Show related menu

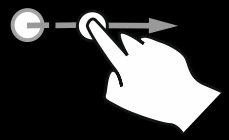

1 finger swipe

- On Whiteboard: Draw Lines
- On Stickynote(the header): Move the stickynote
- On Stickynote(group):Ungroup
- In the command of Move Window: Move the window

2 fingers swipe

- On Document: Next page
- On Whiteboard: Move
- On Stickynote: Move
- On the video area: Switch between panoramic video & video+workspace(in-meeting mode)

2 fingers move inward

- Zoom out
- Scale down the selected objects
- Up to high level of the active chart

2 fingers move outward

- Zoom in
- Enlarge the selected objects
- To detail level of the active chart

5 fingers grab

Switch to overview mode

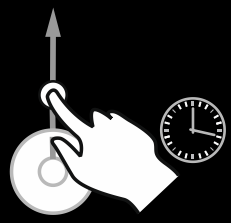

1 finger press & hold, then drag

Share file from/Download file to iPad

Technical Challenge : Example Large Touch Screens

	One user	Multiple users
One device/site	Gesture recognition	User detection Gesture recognition
Multiple devices/sites	Low-latency networking	Low-latency multi-user interactions

Augmented Collaboration Technologies

- Experience simplification!
- Touch/Gesture recognition for large screens
- Remote & Mobile cloud-based collaboration
- Ultra HD video capture, transmission, delay
- Scalable/Multistream video coding and transport
- Networked Video Processing

The evolution to software-based networking: NFV

Network Function Virtualization

The evolution of networks

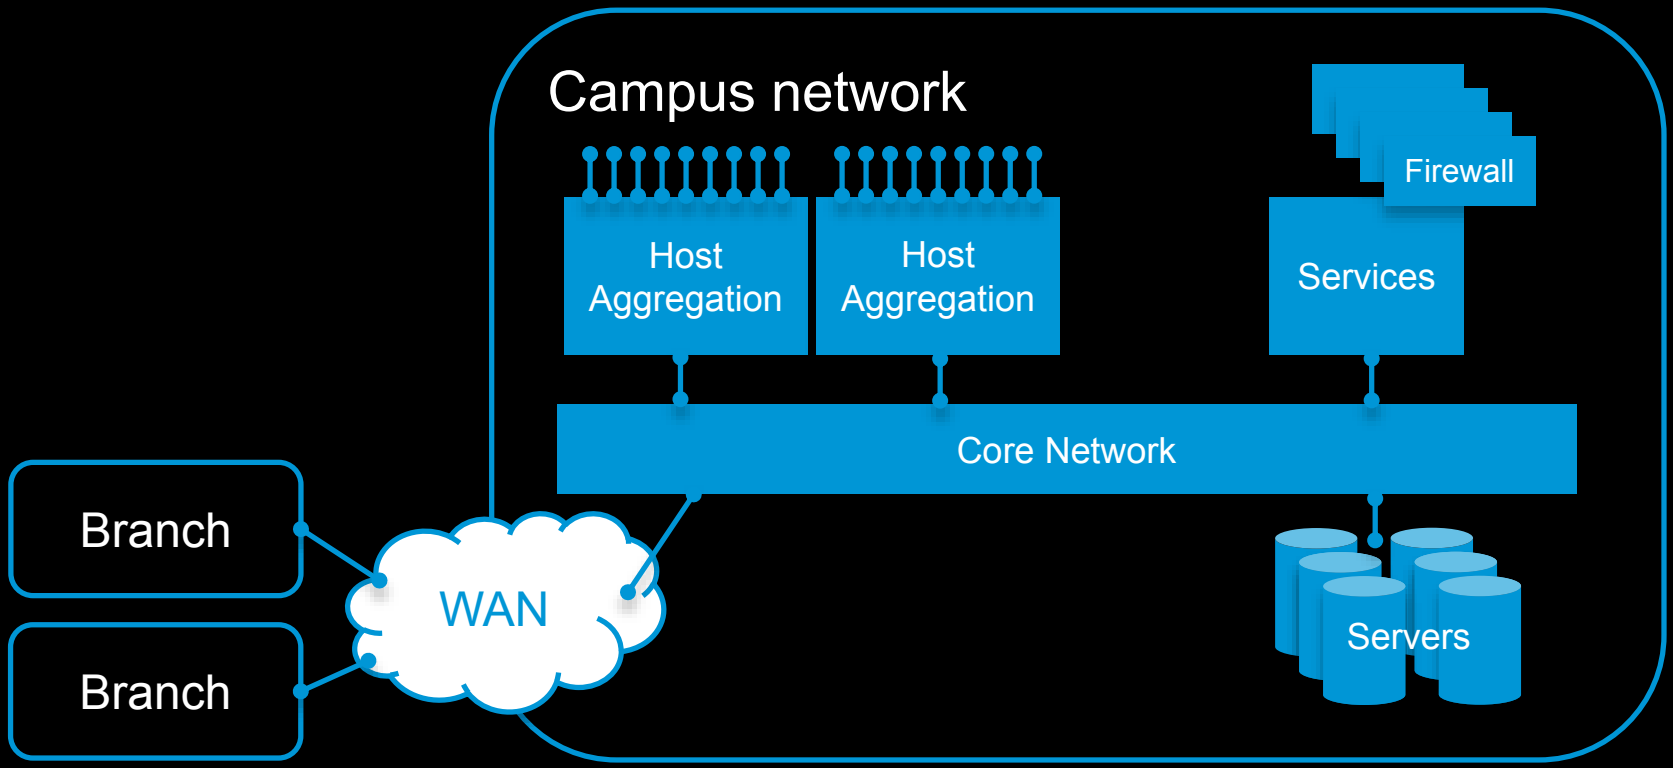

The evolution of networks

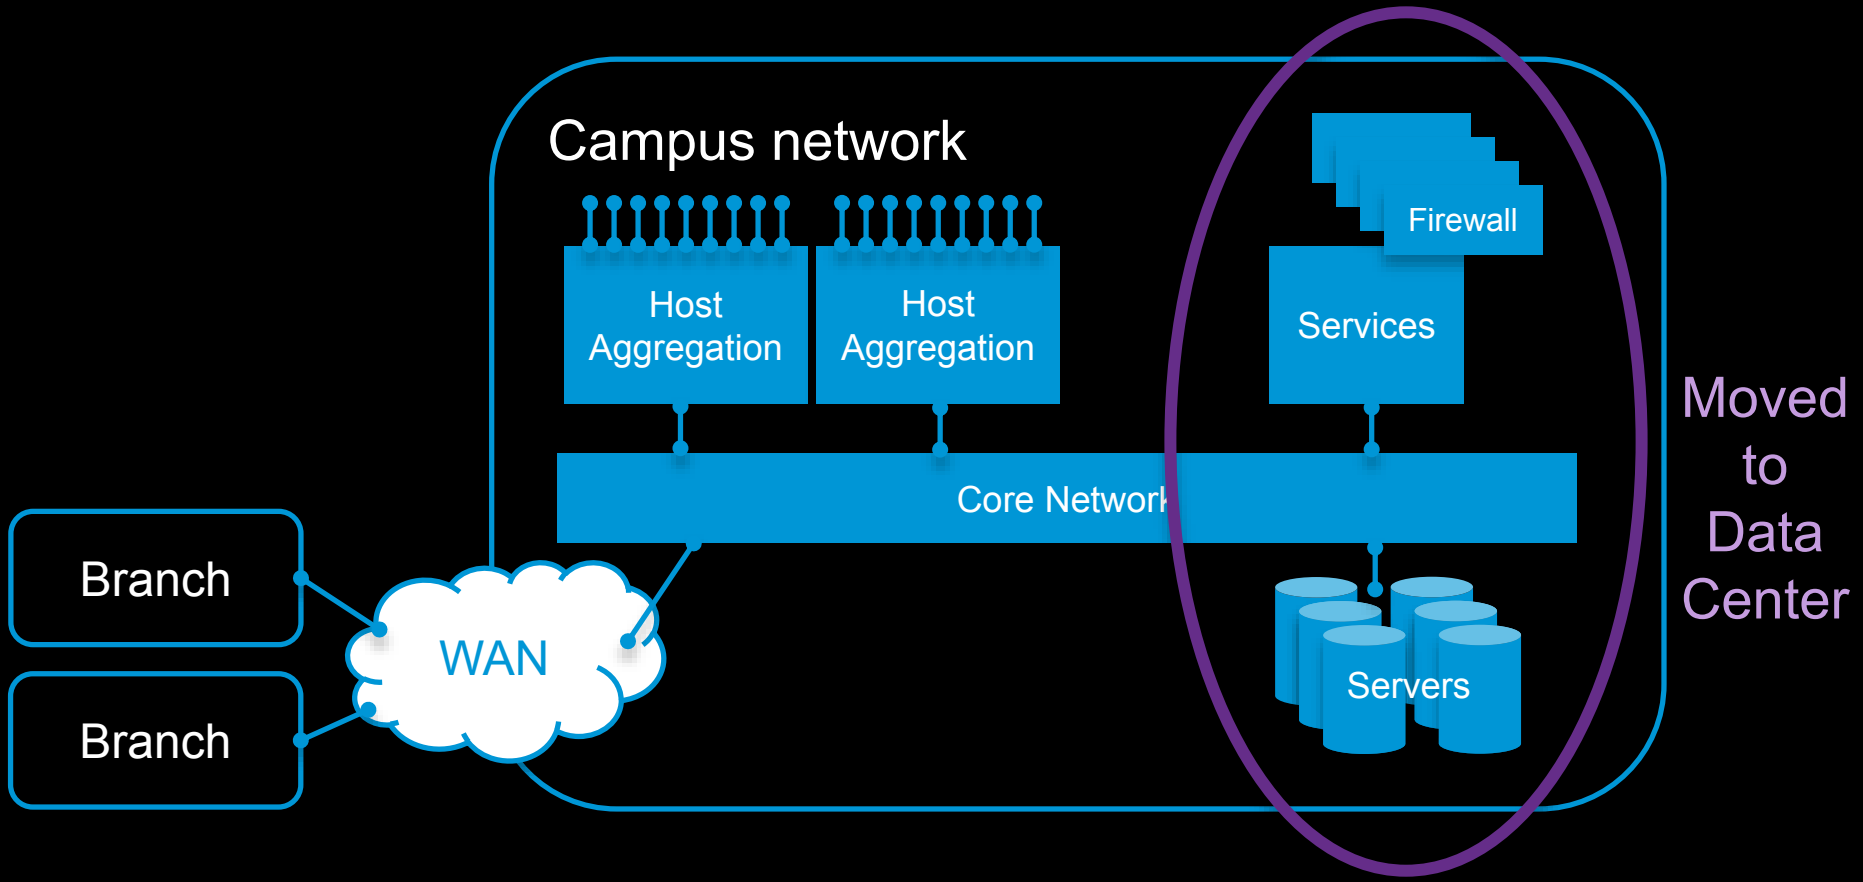

The evolution of networks

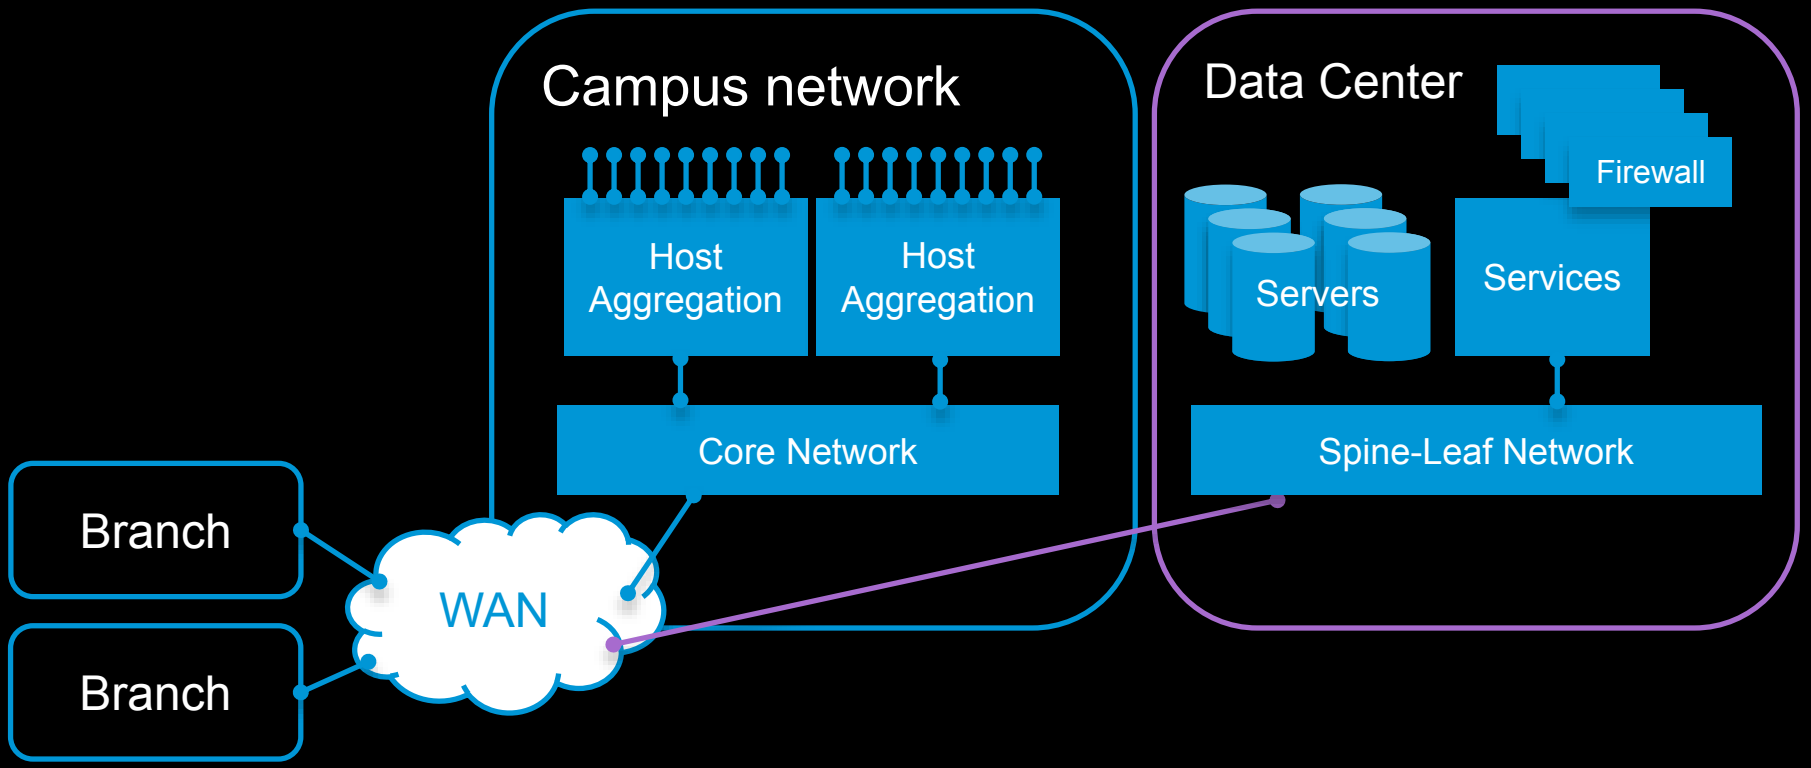

The Right Tool for the Job

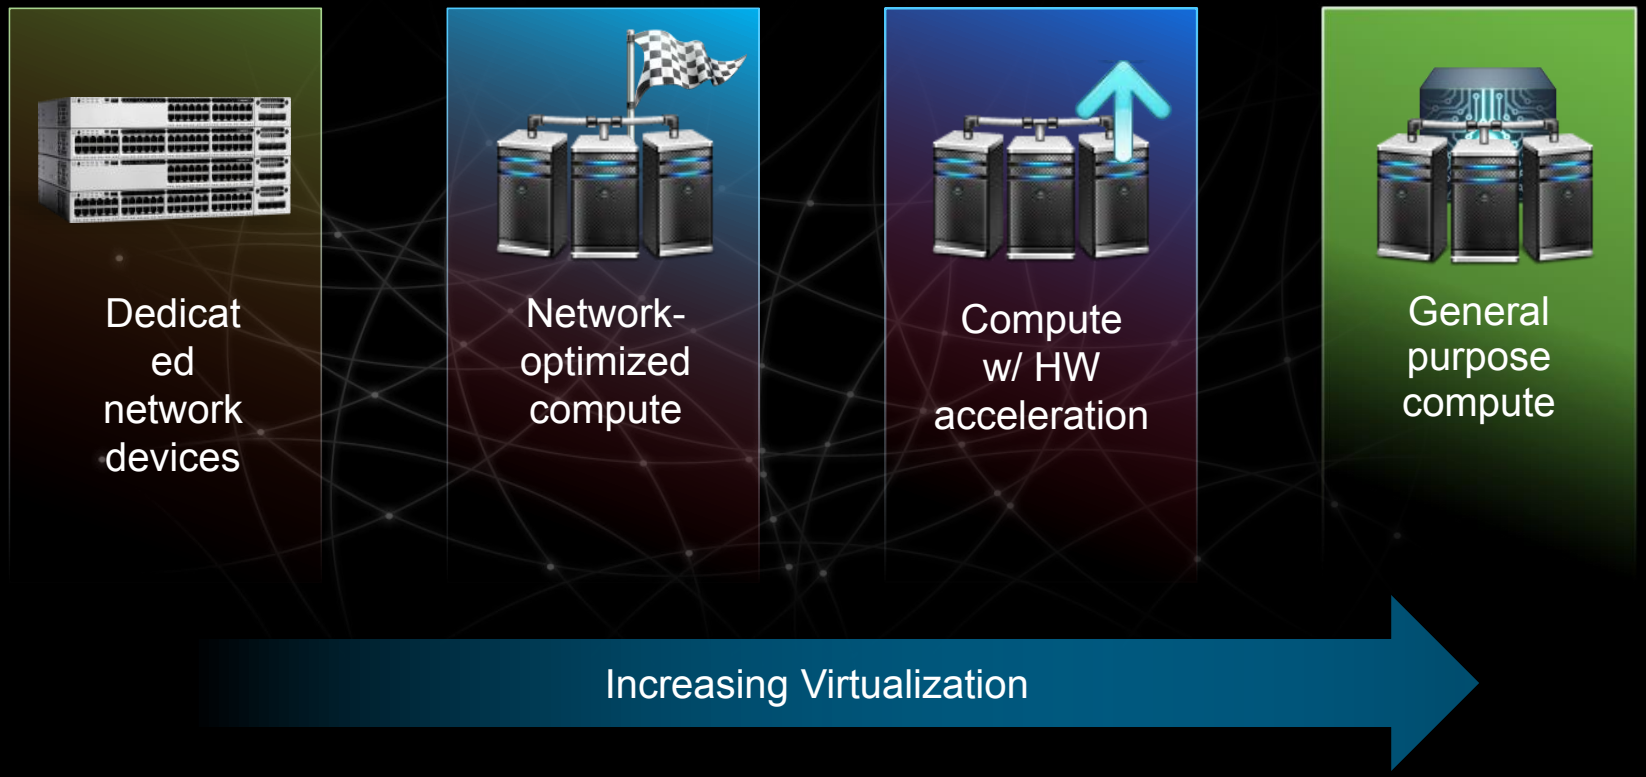

General Purpose Processing

General purpose compute

General Purpose Processing
- Leverages all the platform aspects
  - Power management
  - Legacy interfaces
  - Existing software stacks
  - Platform network drivers
- Can be scaled horizontally using compute virtualization

Networked Optimized Compute

Network-optimized compute

Network Optimized Compute

- Network acceleration for fast packet processing, e.g., Intel Data Plane Development Kit (DPDK) for x86
  - Tight polling loop
  - Turn off power management
- Virtualization not as flexible as general purpose processors in terms of cloud elastic compute, but it provides a balance in flexibility and performance.

Example: NPU ASIC Evolution

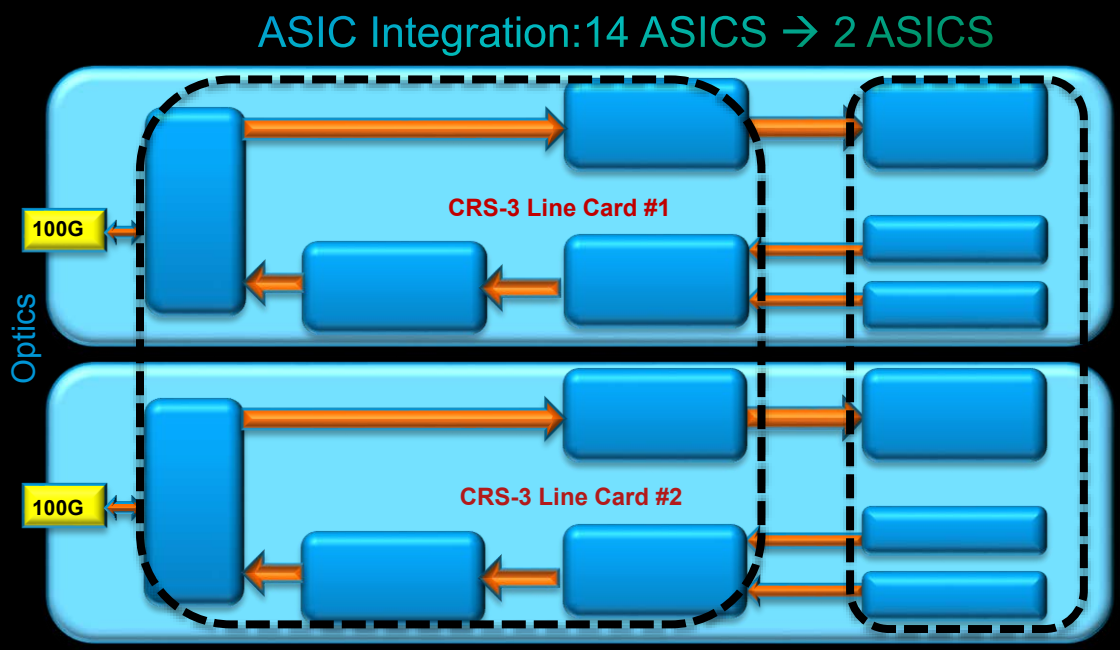

NPU

- 1M+ lines of source code
- 1B+ transistors
- 1000s of signal pins
- > 1 Tbps of I/O BW
- ~100 Watts

Example: NPU ASIC evolution for Cisco nPower X1

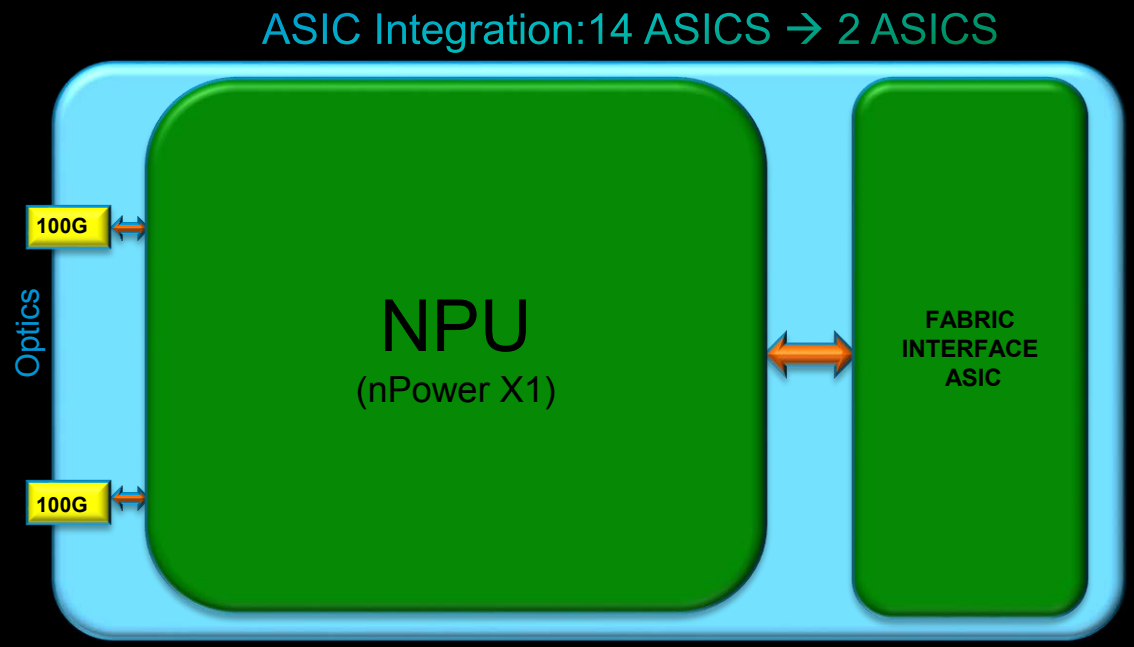

Example: State of the Art NPU

Cisco “nPower X1” NPU

- 4 Billion transistors @ 40nm
- 3199 balls, 55mm package
- 336 multi-threaded processors
  - Fully programmable
- 400Gbps
- 300Mpps
- Custom low-latency memories
- Very power efficient
  - AVS, partial power-down modes

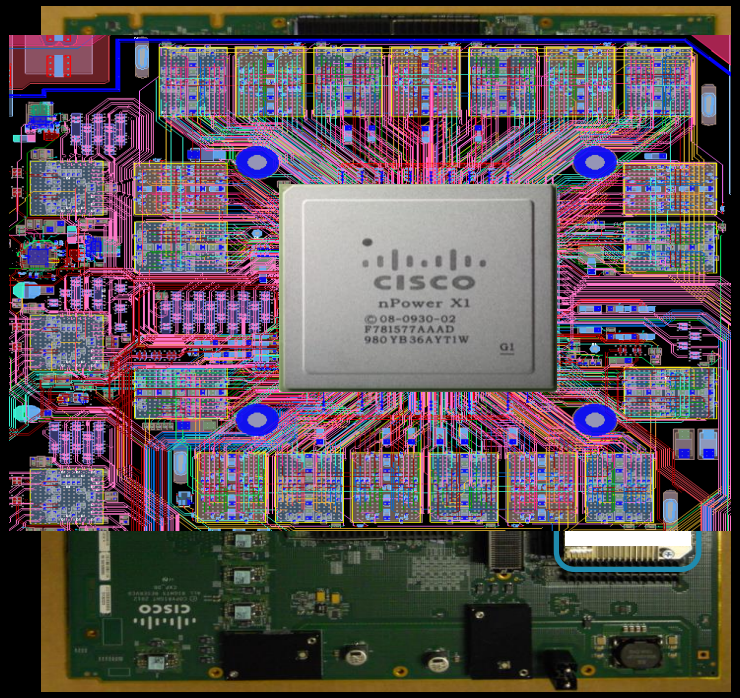

(½ of a) NCS 6000 Linecard

Dedicated Network Devices

High end router technology futures
- Silicon: NPU scaling
- Optics: Silicon Photonics
- System scale: Internal/fabric interconnect
- Packaging: Power/Cooling
- Control Plane: Scale, virtualization

Example NFV Comparison: Power Consumption

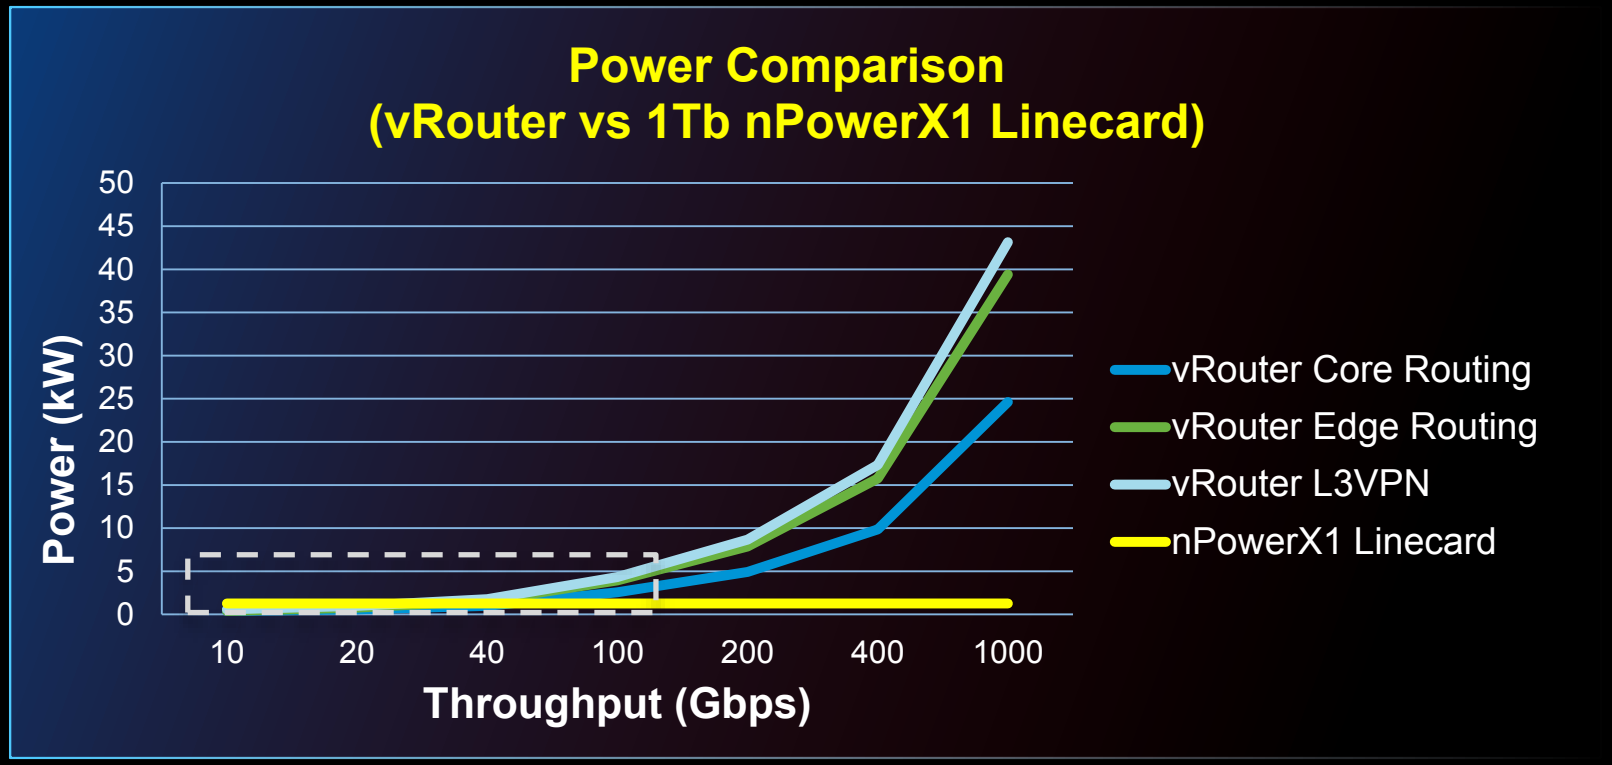

Example NFV Comparison: Power Consumption (zoom in)

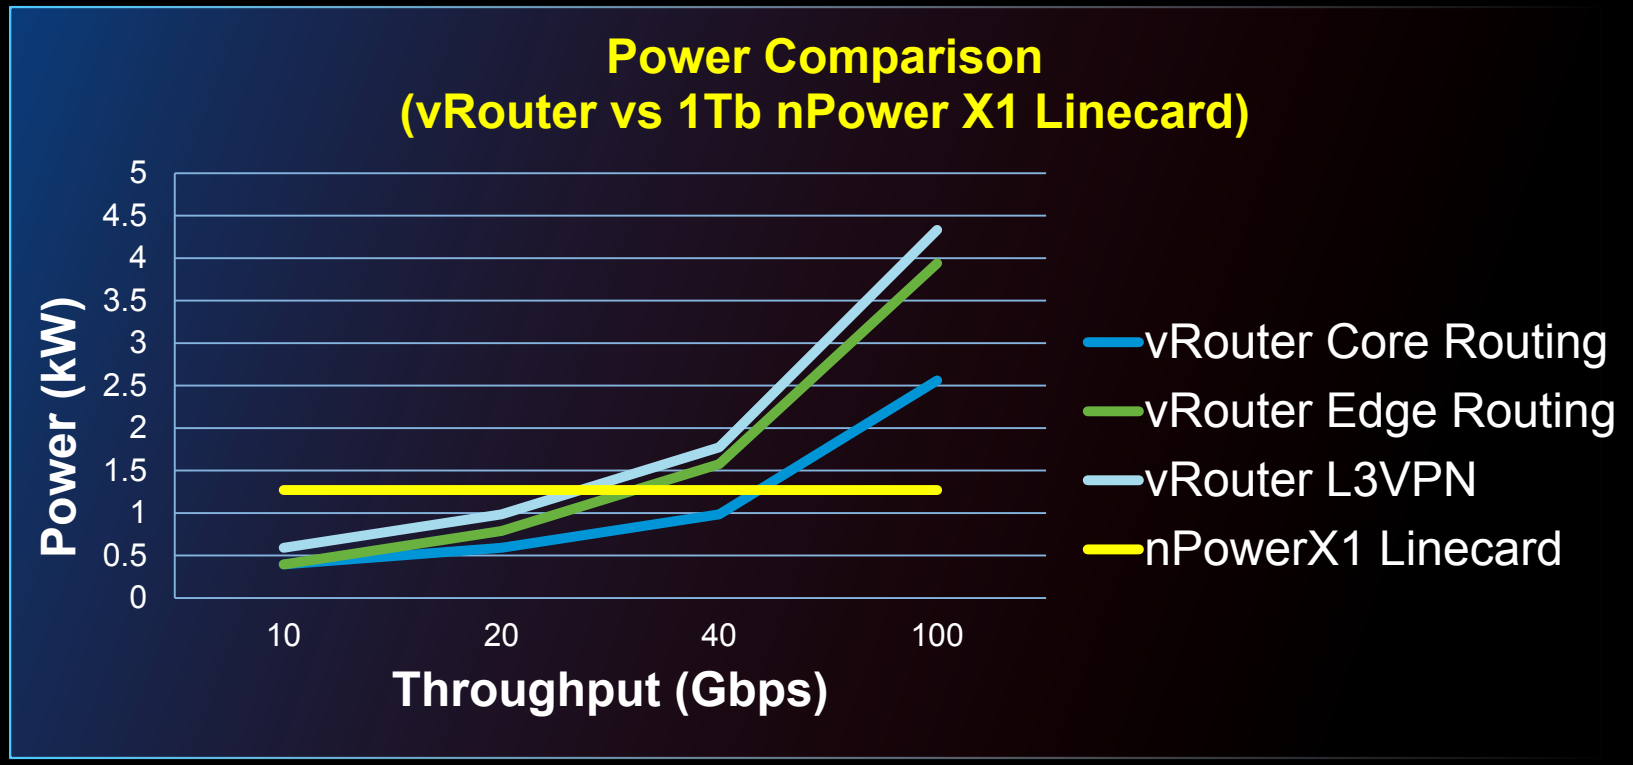

Example NFV Comparison: Number of Blades

Router Throughput	vRouter Blades to Meet Performance			1Tb Physical Router Blades to Meet Performance
	Core Routing	Edge Routing	L3VPN	All Use Cases
100G (10x10G)	13	20	22	1
200G (20x10G)	25	40	44	1
400G (40x10G)	50	80	88	1
1,000G	125	200	219	1
2,000G (20x100)	250	400	437	2
8,000G (80x100)	1000	1600	1746	8

Advantages of pure software NFV

- Flexible & faster deployment of new software services
- Removes the number of specialized physical boxes
- Allows elastic use of compute resources
- Reduces operational costs by oversubscribing hardware beyond what they can do on individual nodes
- Reduces risk of stranded capex if a specific project or service doesn't take off

Network solutions: Requirements and today's approaches

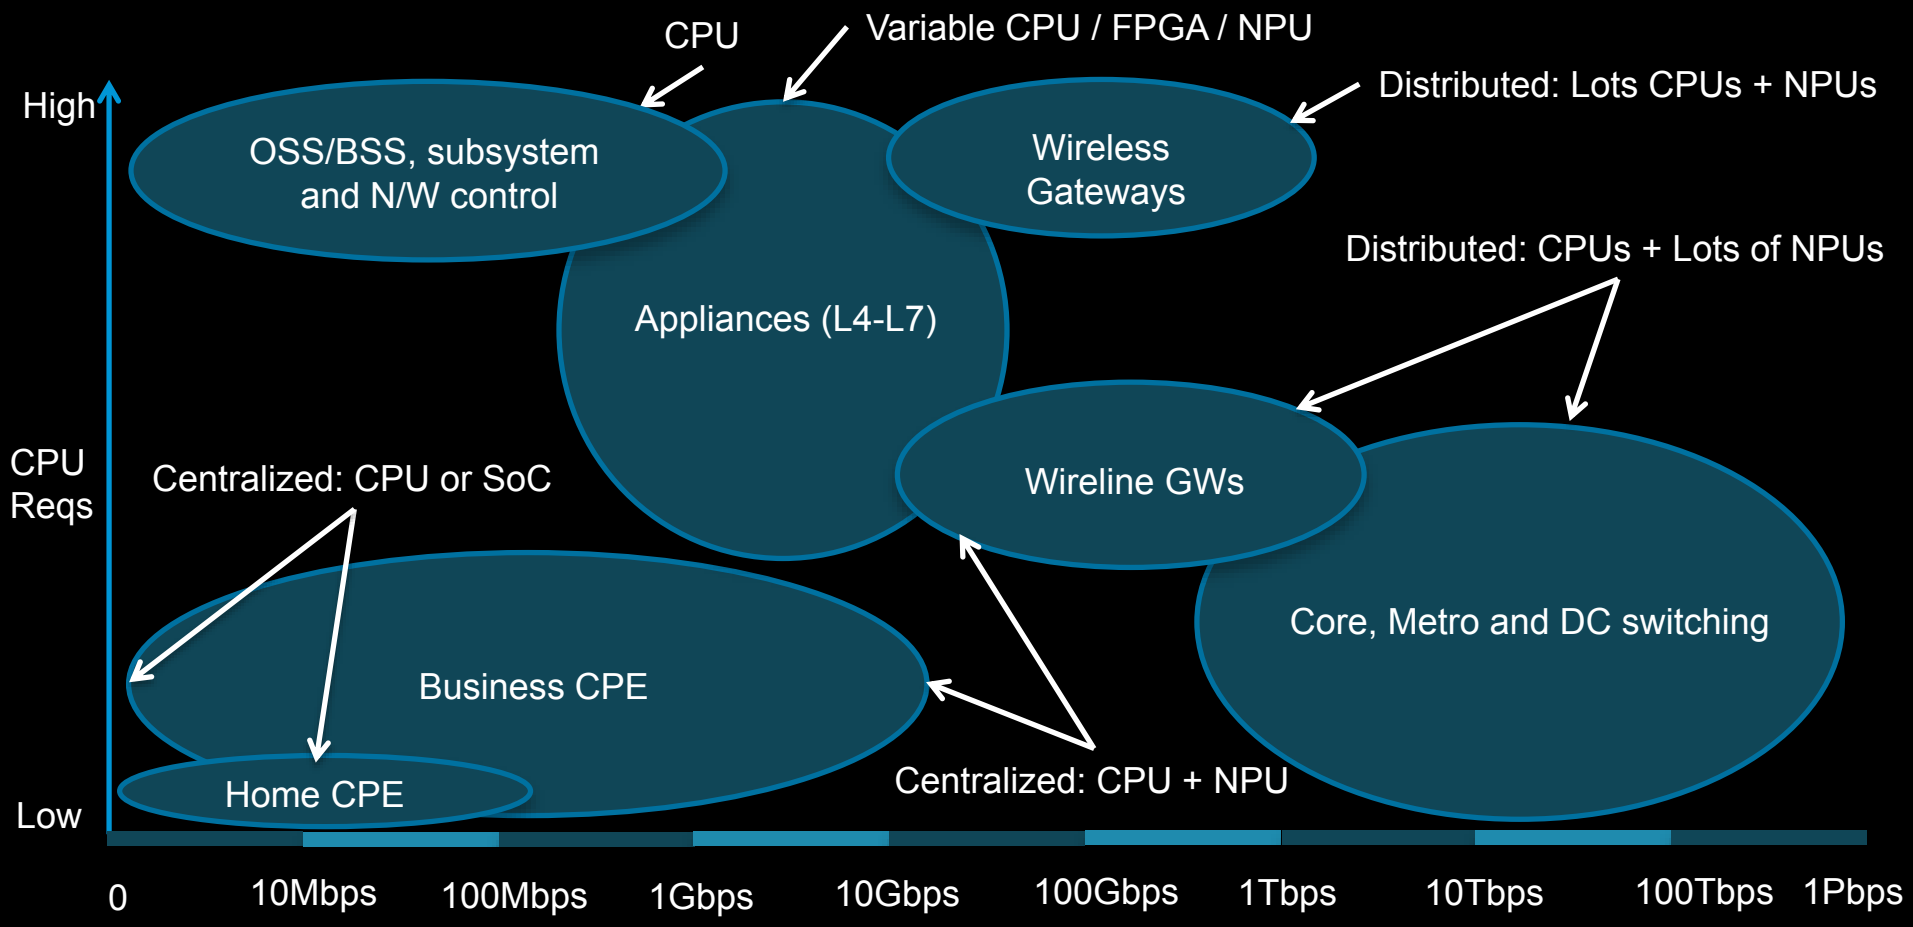

NFV Conclusion

- Dedicated network hardware is more power efficient than general purpose compute and achieves orders of magnitude improvement for packet processing and optimized network functions.
- General purpose compute helps when deploying new applications and services that do not have optimized network functions. Over time, the most highly used services may get baked into the hardware-optimized implementations on devices.
- *Use the right tool for the right job. Architect for the future.*

NeXt Experiences for Software-Defined Networking (SDN)

The Network Operator Experience

Challenges of operating networks at scale

Complex and Unwieldy	Inability to Deploy New Services	Demanding Highly skilled Staff	Difficult to Provision Applications	Difficult to Troubleshoot	Fragile Security Infrastructure

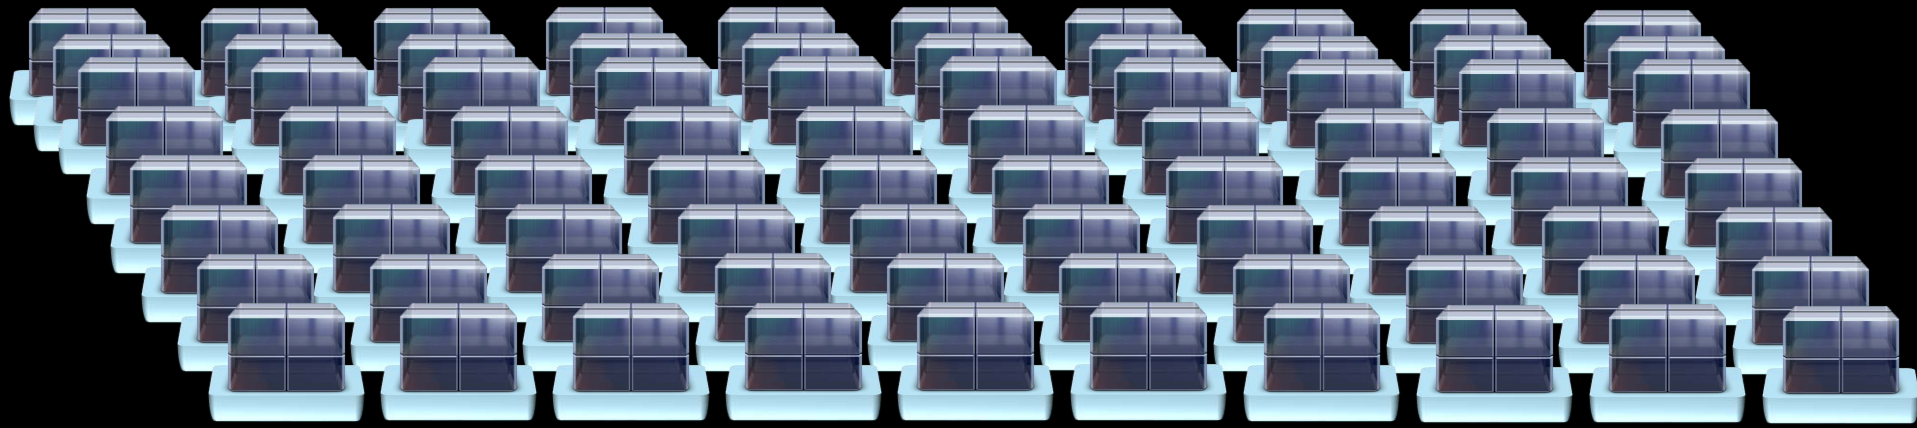

100s → 1,000s → 10,000s → ????

The new architecture of networks

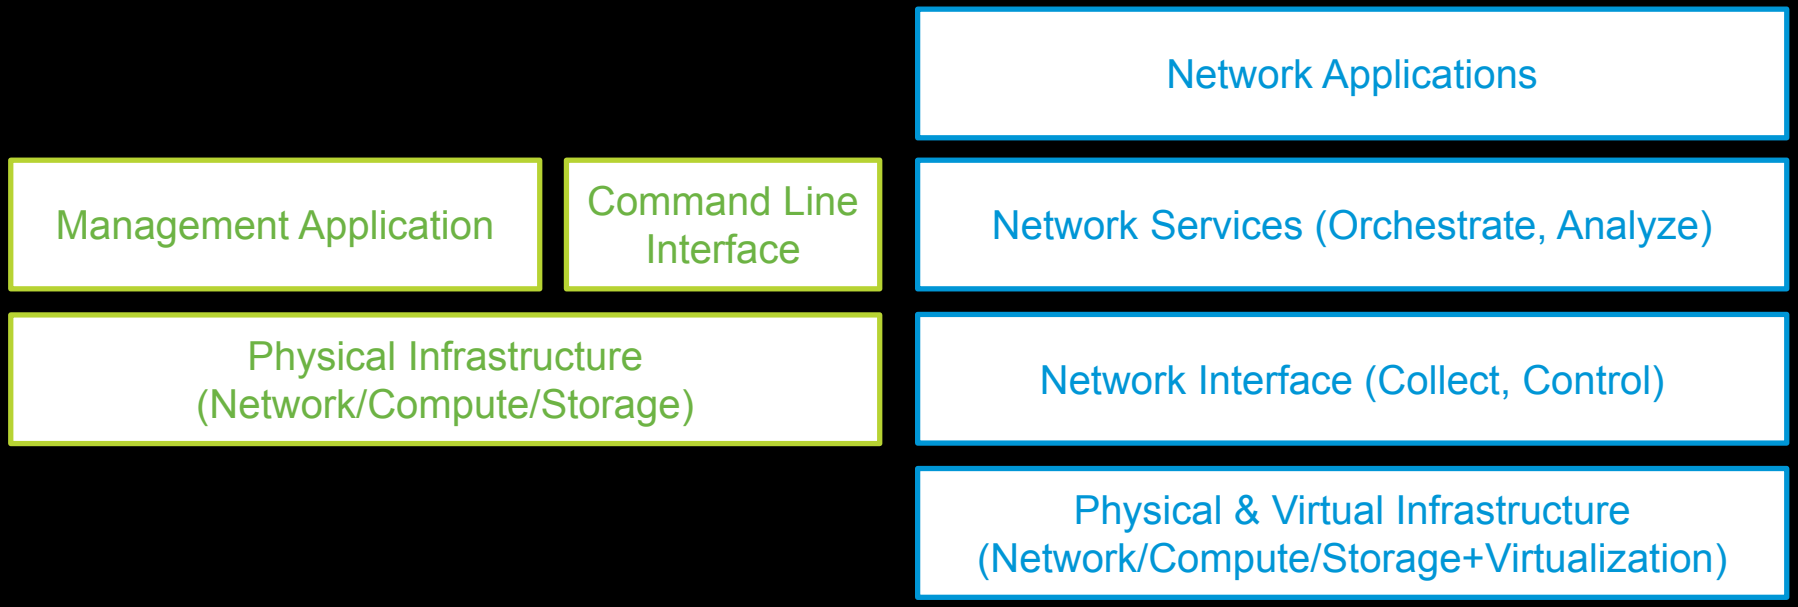

Cisco ONE SDN/NFV: A Modular Solution Framework

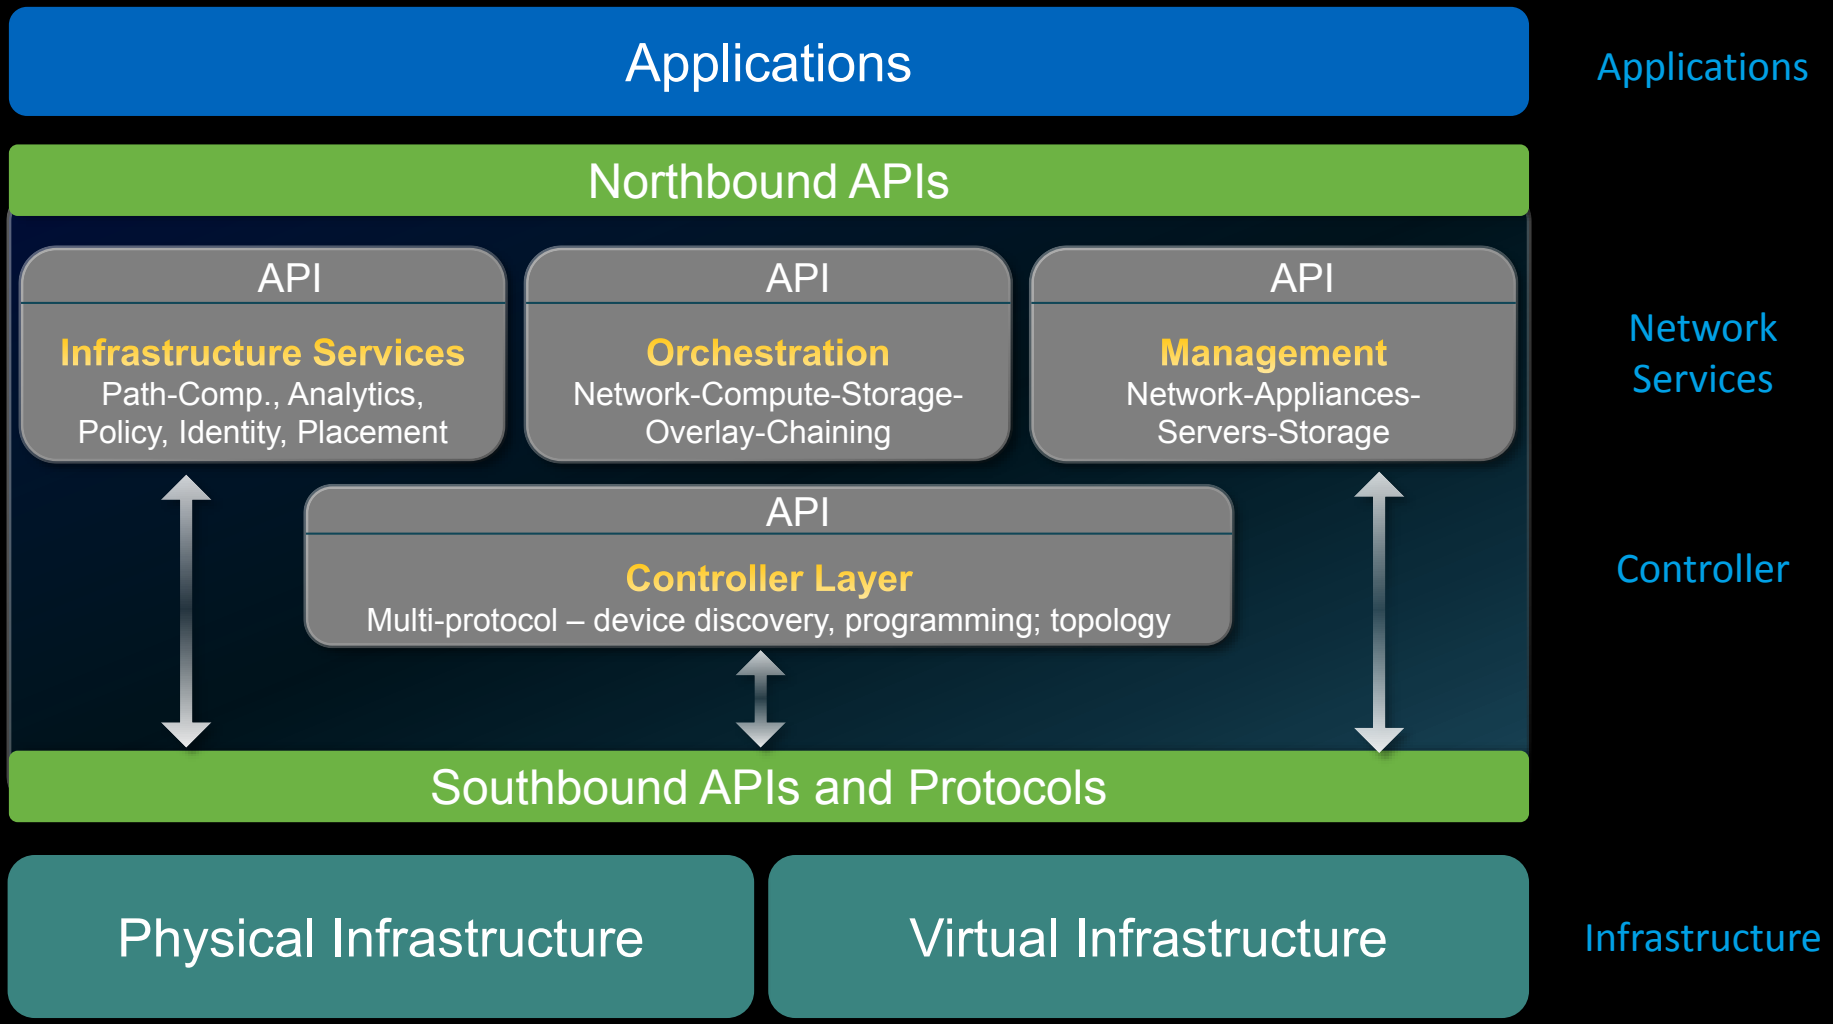

SDN: Domain Controller with Workflow Orchestration

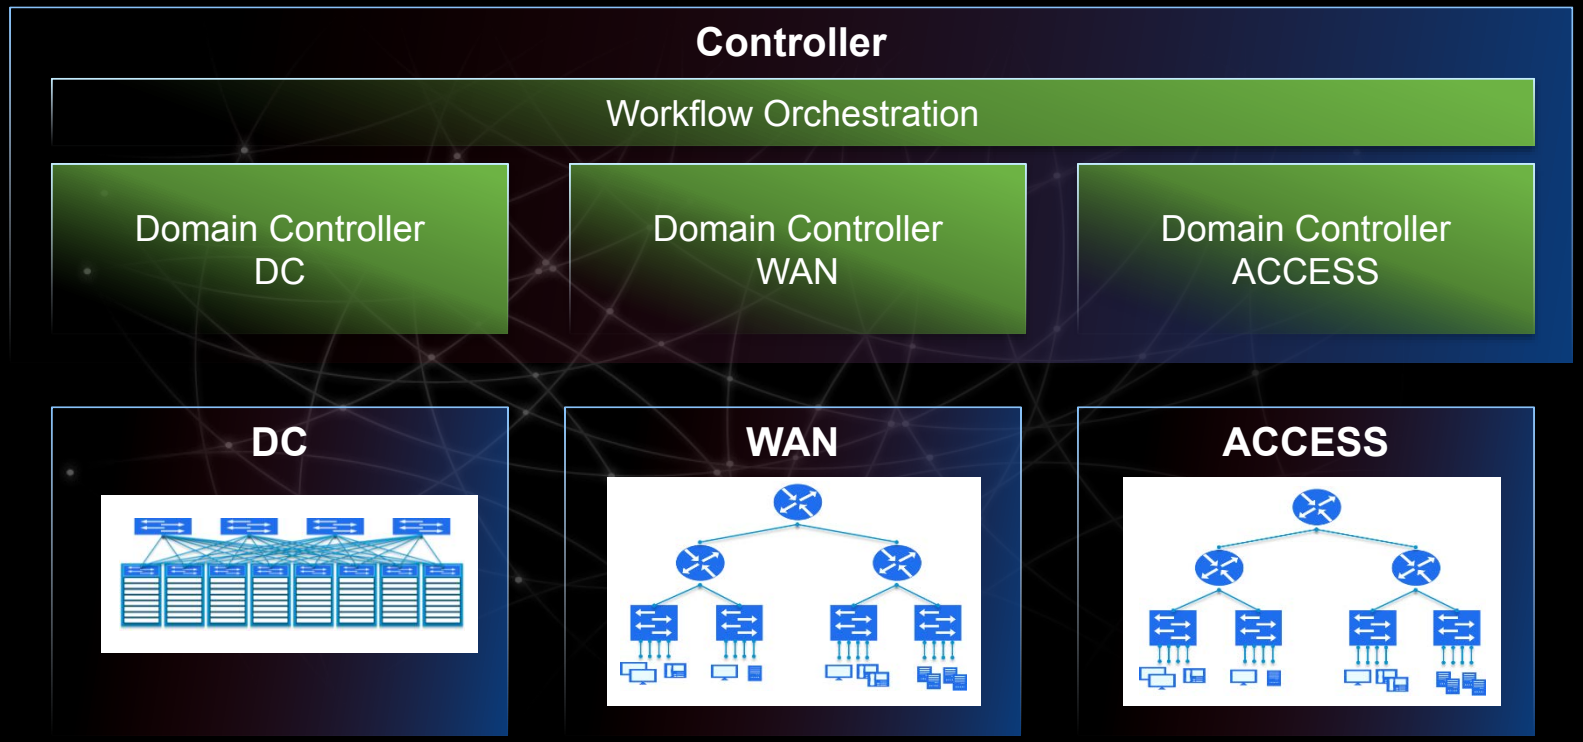

Networked Experience Leap

From CLI to NeXt

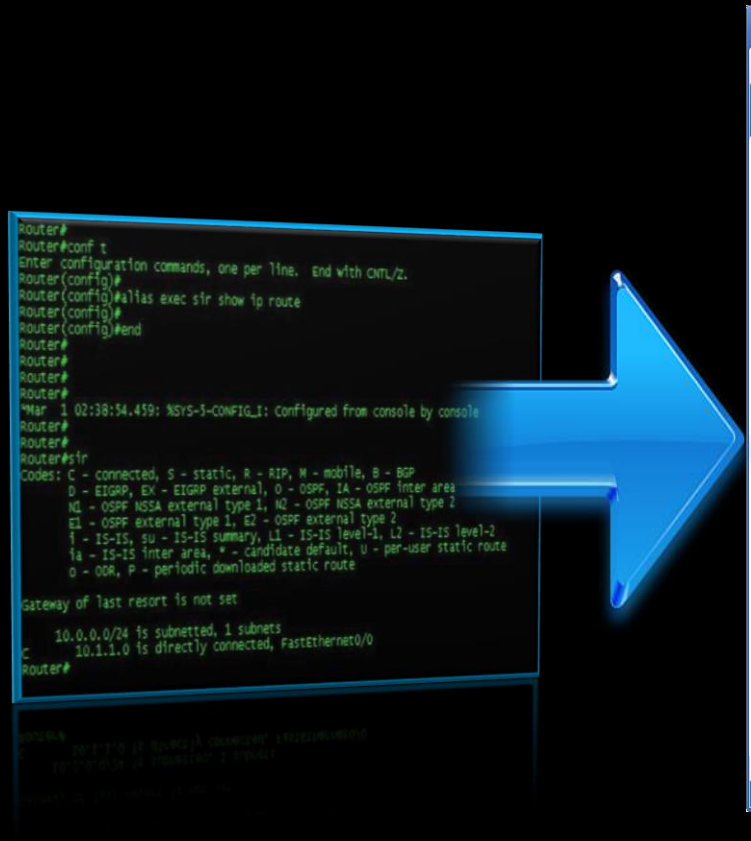

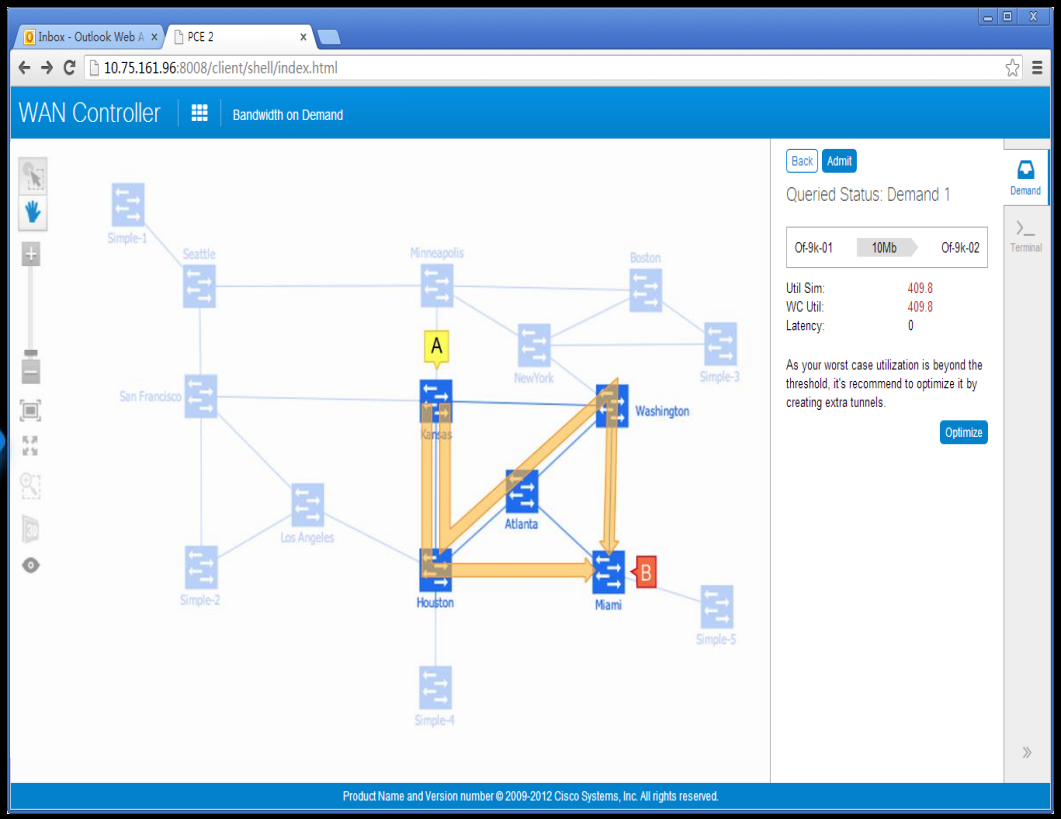

Enterprise SDN: Example use cases for Cisco APIC EM

Security Automation

Network-Wide Rapid Threat Detection and Mitigation (Sourcefire)

ACL Management Automation

QoS Provisioning

Easy QoS

Follow Me QoS

Compliance Assurance

IWAN: Path Optimization

Automated Performance Routing (PfR) Configuration

Automated WAN Policy Compliance Assurance

SDN Example: Topology visualization in Cisco APIC EM

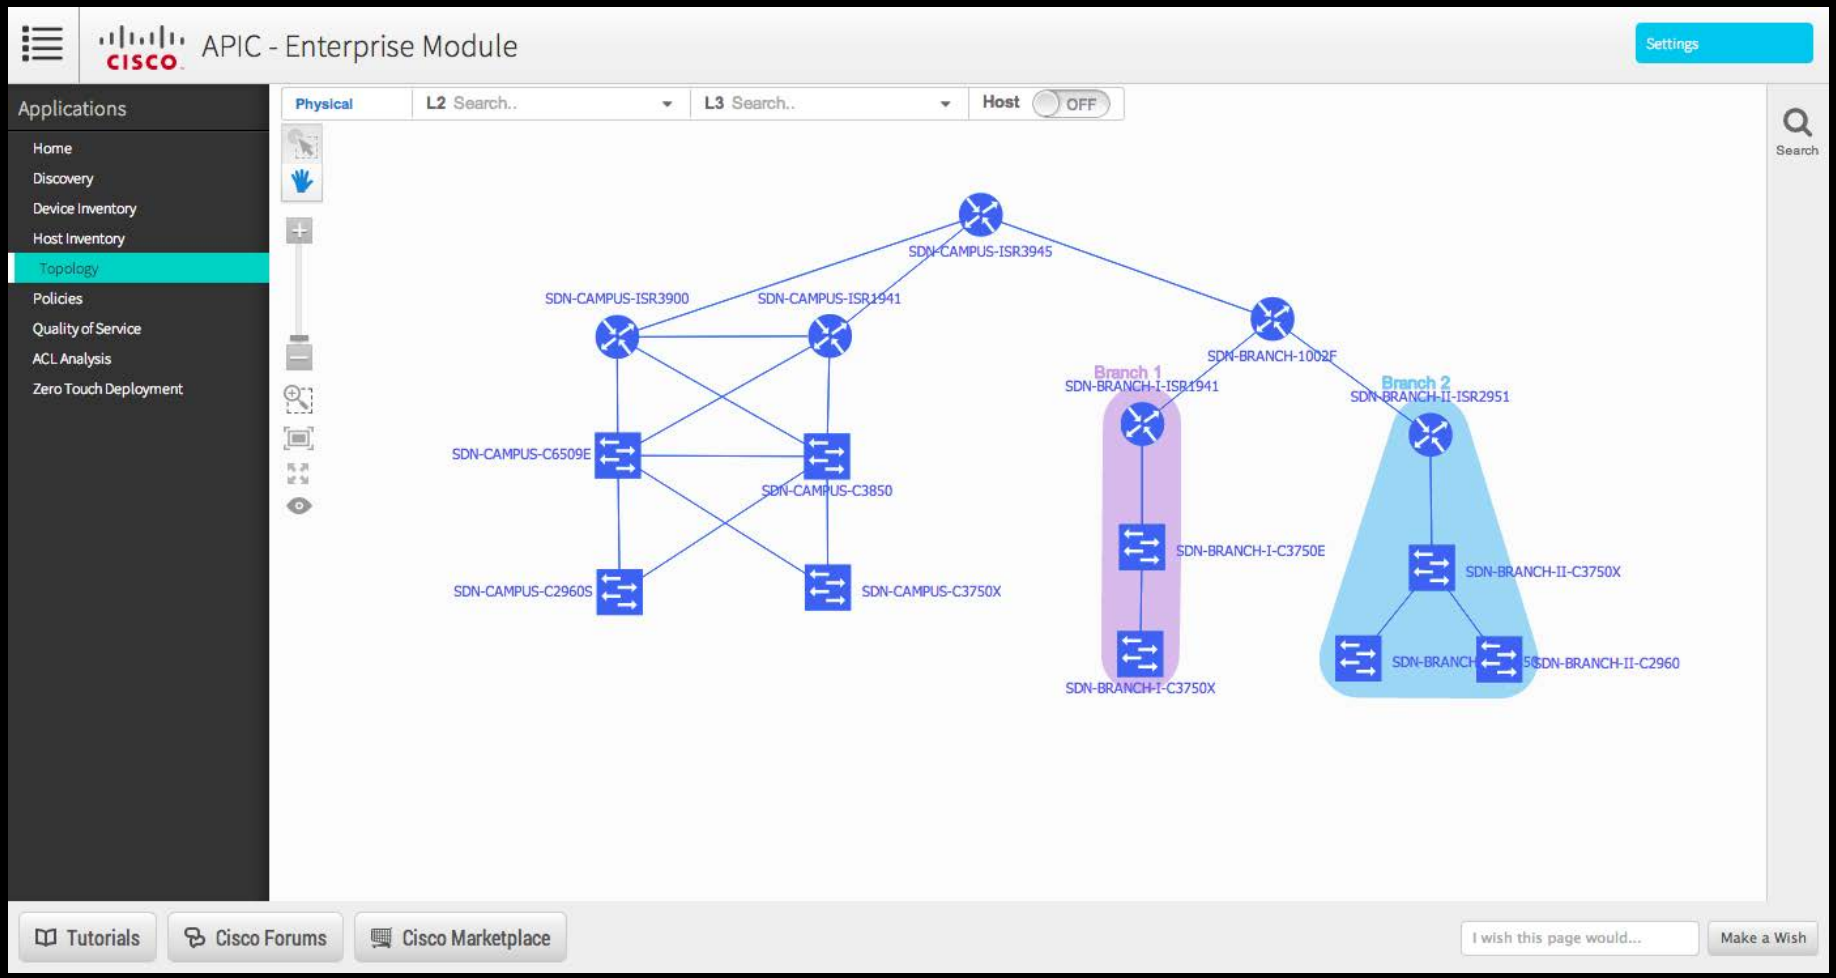

SDN Example: Access Control List (ACL) Management in Cisco APIC EM

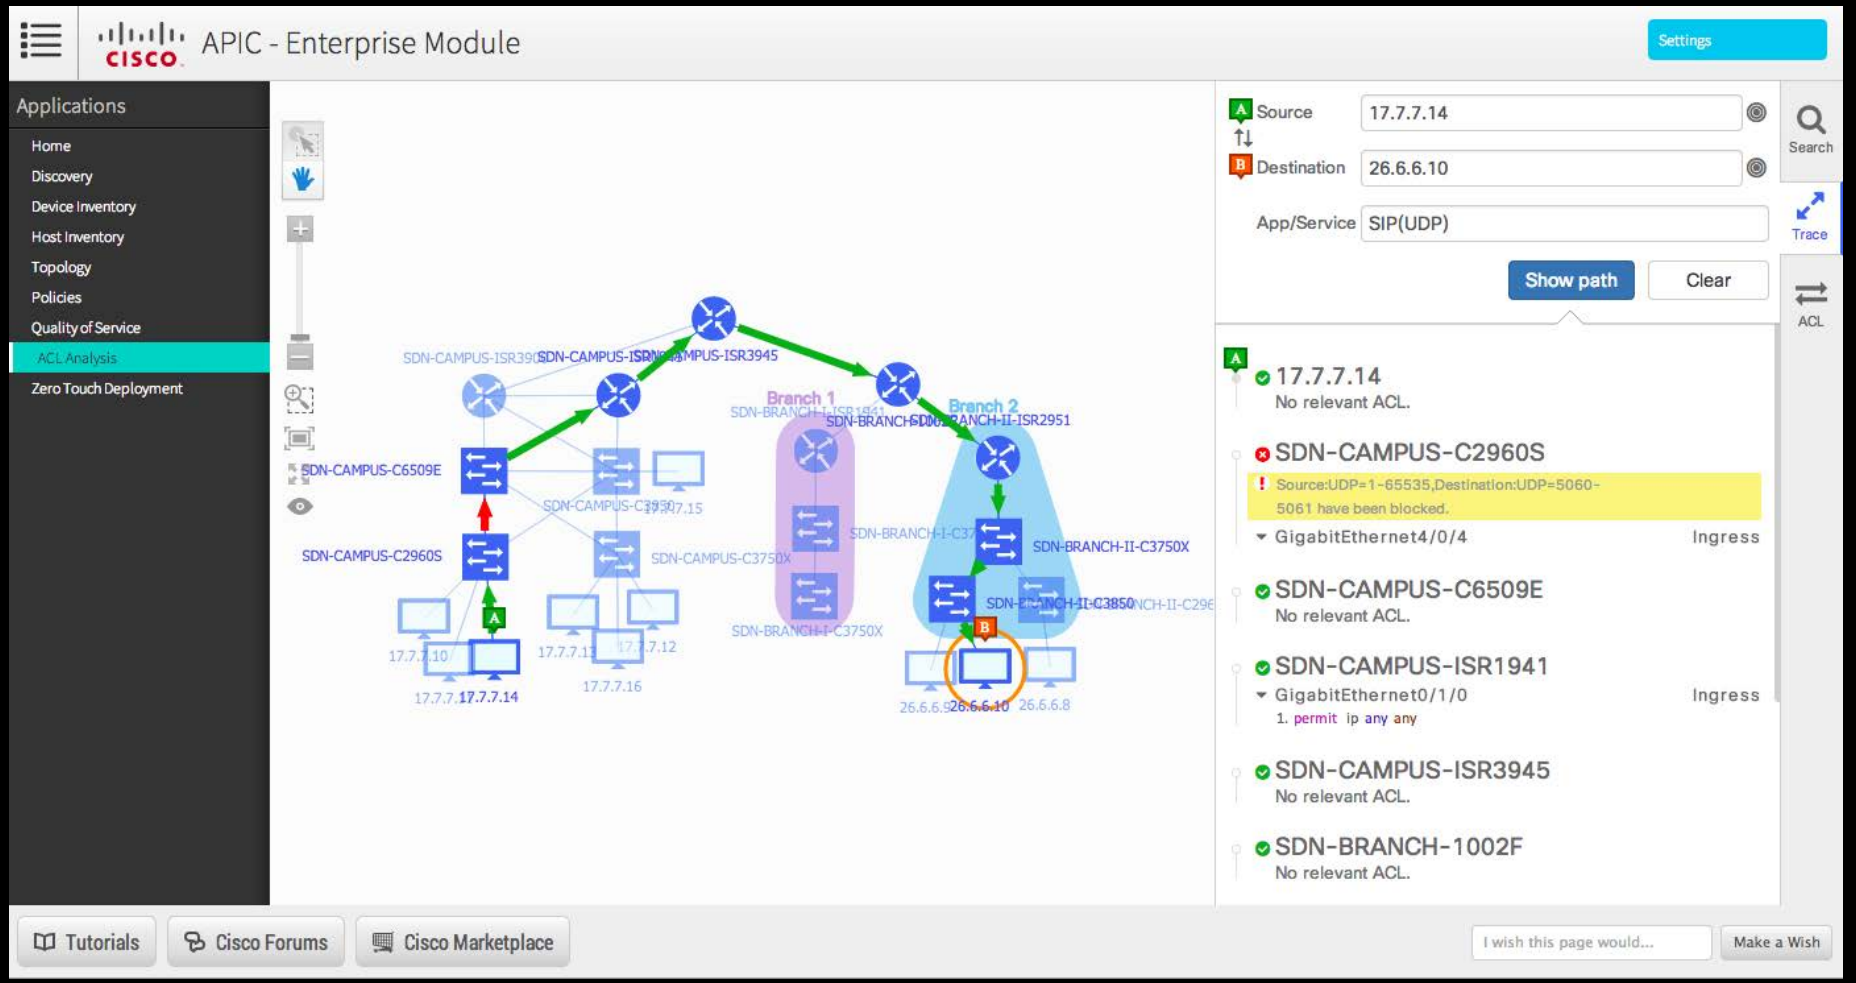

SDN Example: Cisco SP WAN Orchestration

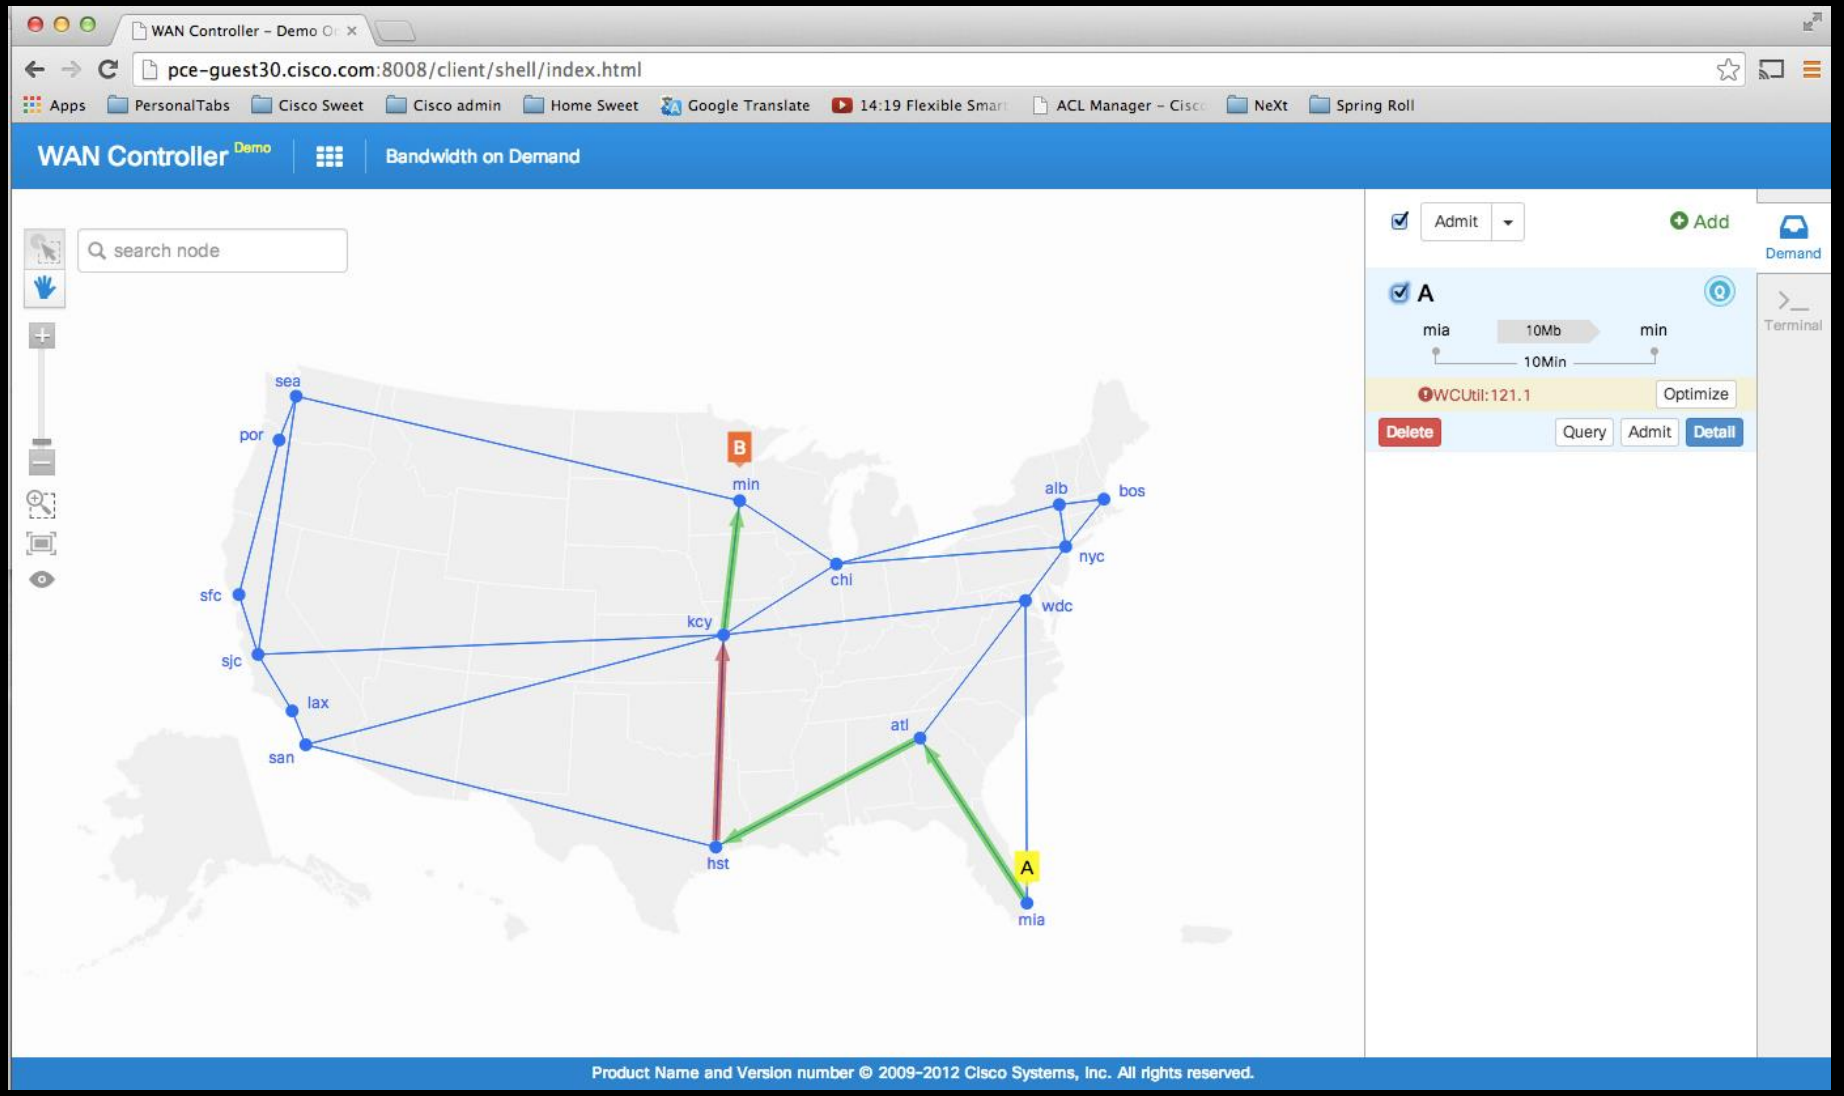

Example application : Intelligent WAN orchestration

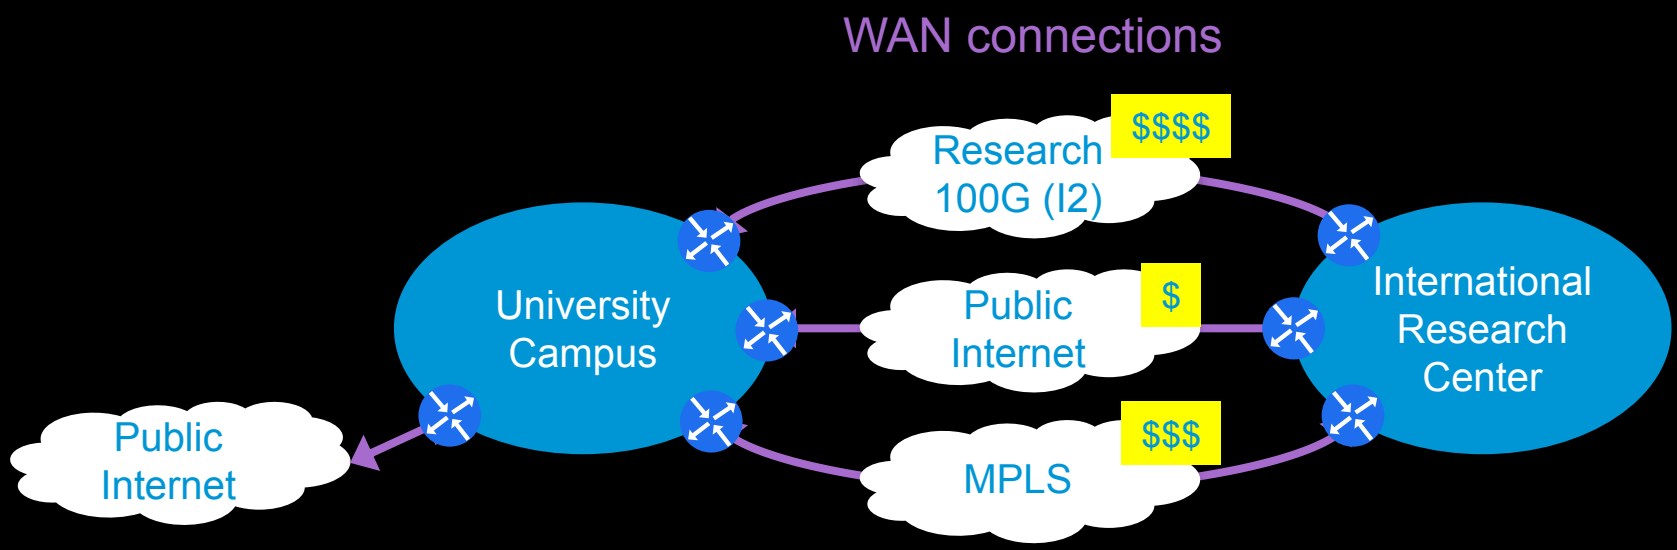

Example application : Intelligent WAN orchestration

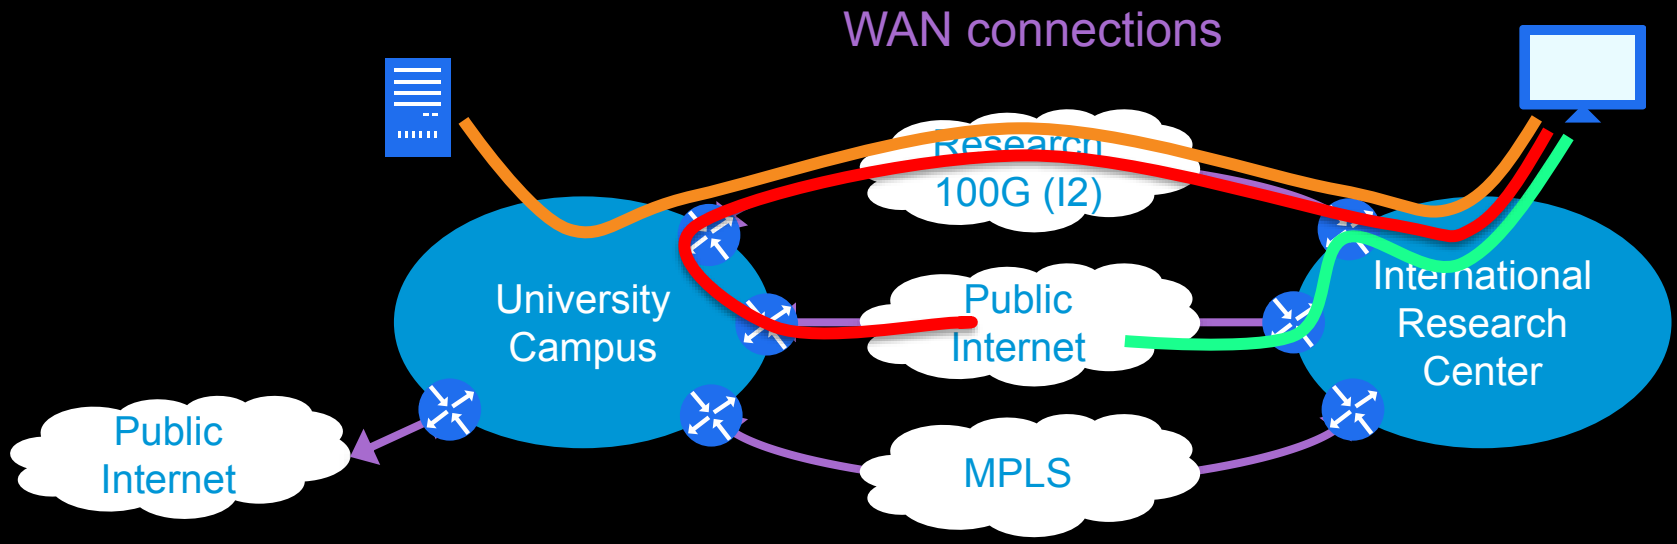

Cisco – APIC Enterprise Module: Automatic Threat Detection and Mitigation

Network Wide Security Deployed Rapidly

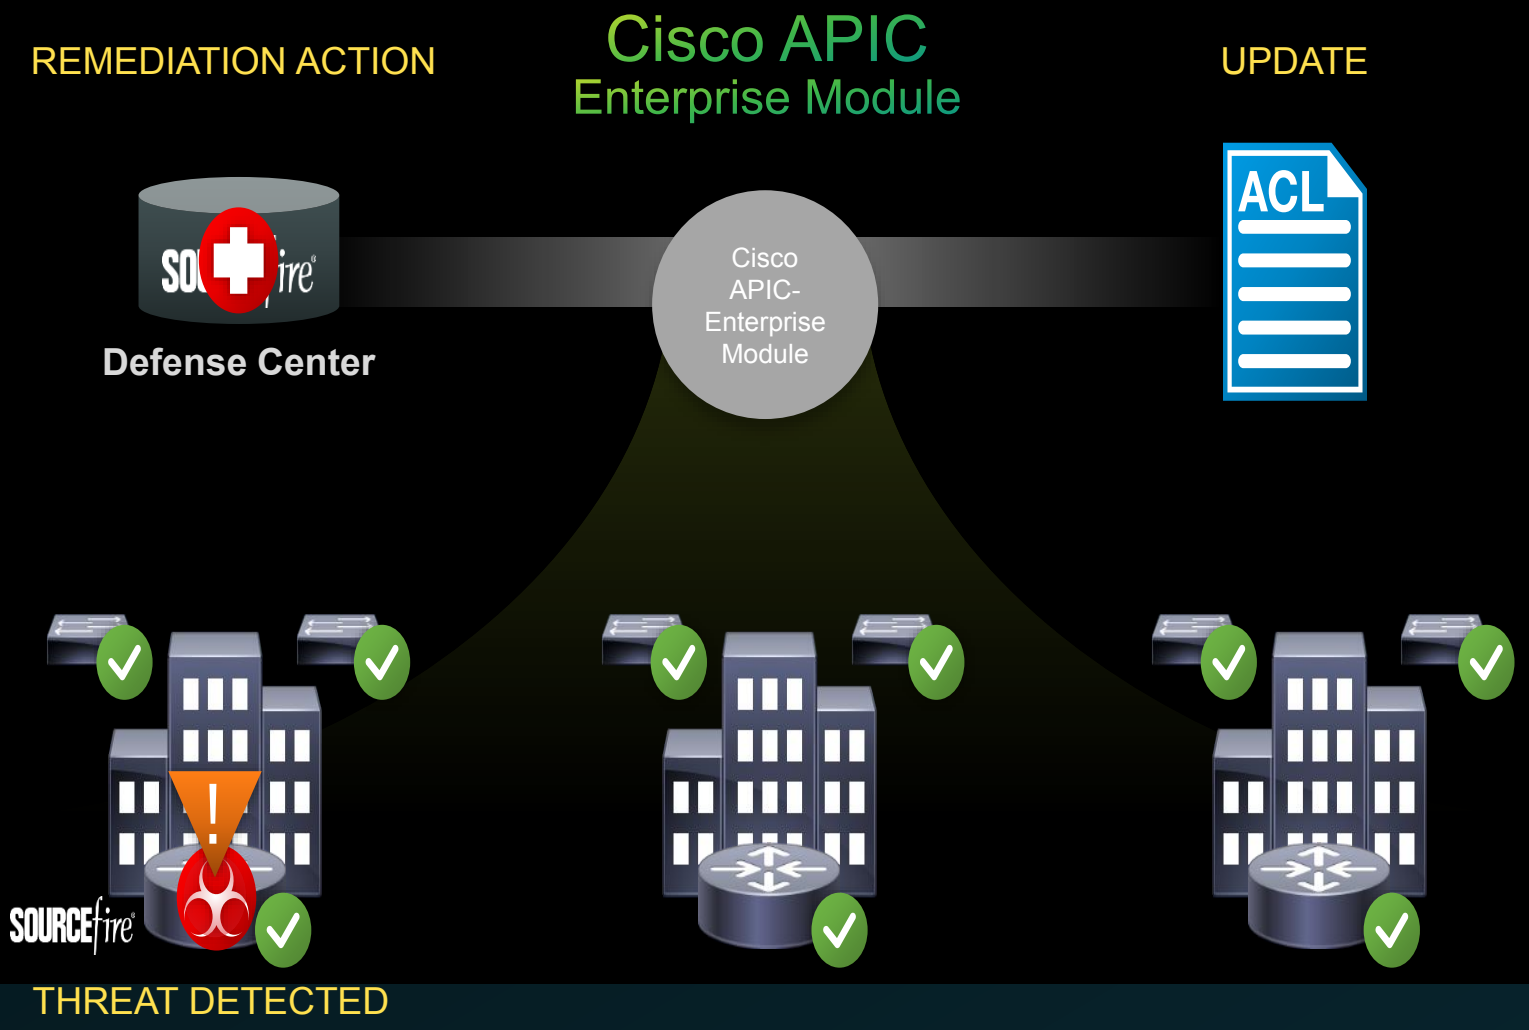

Conclusions

- Networking is moving towards network programmability with SDN and NFV.
- Applications can be tied to the network, enabling compelling experiences like Augmented Collaboration.
- Innovations in ASICs and dedicated routers continue to achieve dramatically increasing performance.
- NFV raises the possibility of network functions on general purpose processors (CPUs).
- There is a place for hardware and software- Use the right tool for the right job.
- SDN provides network-wide visibility and control.
- SDN/NFV allows network services to be added to the network and accessed through programmatic interfaces (APIs).
- SDN performs best when using all the network element capabilities. If dedicated network hardware is available, use it.
- Big opportunity lies in optimizing workflows with awareness of network element capabilities. Power-aware SDN/NFV.

Thank you.